



10-CV-01385-EXH

EXHIBIT 1

About Google!

Links

- [Research Papers about Google and the WebBase](#)
- [Pictures and stats for the Stanford Google Hardware](#)

Credits

- Current Development: [Sergey Brin](#), [Larry Page](#), and [Craig Silverstein](#)
- Design and Implementation Assistance: [Scott Hassan](#) and [Alan Steremberg](#)
- Faculty Guidance: [Hector Garcia-Molina](#), [Rajeev Motwani](#), [Jeffrey D. Ullman](#), and [Terry Winograd](#)
- Research Funding: [NSF](#), [NASA](#), [DARPA](#) and [Interval Research](#)
- Equipment Donations: [IBM](#), [Intel](#), and [Sun](#)
- Equipment Consulting: [Penguin Computing](#)
- Software: [GNU](#), [Linux](#), [Python](#), [Parasoft](#) (debugging), and [Gimp](#) (logo design)
- Collaborating Groups in the [Computer Science Department](#) at [Stanford University](#): [The Digital Libraries Project](#), [The Project on People Computers and Design](#), [The Database Group](#), [The Stanford InfoLab](#), [The MIDAS Data Mining Group](#), and [The Theory Division](#)
- Outside Collaborators: [Interval Research Corporation](#) and the [IBM Almaden Research Center](#)
- Technical Assistance: [The Computer Science Department's Computer Facilities Group](#), [Stanford's Distributed Computing](#) and [Intra-Networking Systems Group](#)

Last modified: Sun Sep 27 20:26:46 PDT 1998

EXHIBIT 2

US006263507B1

(12) **United States Patent**
Ahmad et al.

(10) **Patent No.:** **US 6,263,507 B1**
 (45) **Date of Patent:** **Jul. 17, 2001**

(54) **BROWSER FOR USE IN NAVIGATING A BODY OF INFORMATION, WITH PARTICULAR AFFILIATION TO BROWSING INFORMATION REPRESENTED BY AUDIOVISUAL DATA**

(75) **Inventors:** Subutai Ahmad; Neal A. Bhadhauskar, both of Palo Alto; Steve B. Canning, Mountain View, all of CA (US); Emmanuel E. Farber, New York, NY (US); Paul A. Freilinger, San Mateo, CA (US); Christopher D. Horner, Redmond, WA (US); Philippe P. Pierrot, Palo Alto, CA (US); Brygg A. Ullmer, Cambridge, MA (US)

(73) **Assignee:** Interval Research Corporation, Palo Alto, CA (US)

(*) **Notice:** Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) **Appl. No.:** 08/761,030

(22) **Filed:** Dec. 5, 1996

(51) **Int. Cl.⁷** H04N 7/173; H04N 5/445; G06F 15/16

(52) **U.S. Cl.** 725/134; 709/217; 725/38; 725/100; 725/110; 725/133

(58) **Field of Search** 345/327; 709/217-219; 348/6, 7, 10, 12, 13, 8; 455/3.1, 4.1, 4.2, 5.1, 6.1, 6.2, 6.3; 725/37-41, 43, 86, 87, 100, 109, 110, 131, 133, 134, 139, 141, 142, 151, 153

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,406,626 4/1995 Ryan 380/9
 5,428,774 6/1995 Takahashi et al. 395/600

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

WO 96/12240 4/1996 (WO) .

OTHER PUBLICATIONS

Elliot, E., "Multiple Views of Digital Video", MIT Media Laboratory, Interactive Cinema Group, Mar. 23, 1992.

(List continued on next page.)

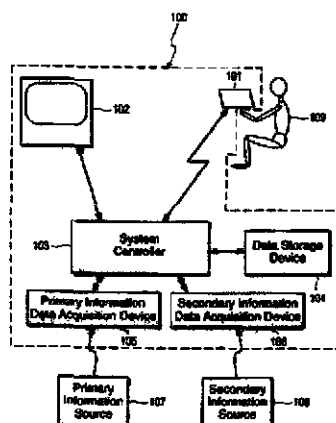
Primary Examiner—John W. Miller

(74) *Attorney, Agent, or Firm*—David R. Graham

(57) **ABSTRACT**

The invention facilitates and enhances review of a body of information (that can be represented by a set of audio data, video data, text data or some combination of the three), enabling the body of information to be quickly reviewed to obtain an overview of the content of the body of information and allowing flexibility in the manner in which the body of information is reviewed. In a particular application of the invention, the content of audiovisual news programs is acquired from a first set of one or more information sources (e.g., television news programs) and text news stories are acquired from a second set of one or more information sources (e.g., on-line news services or news wire services). In such a particular application, the invention can enable the user to access the news stories of audiovisual news programs in a random manner so that the user can move quickly among news stories or news programs. The invention can also enable the user to quickly locate news stories pertaining to a particular subject. Additionally, when the user is observing a particular news story in a news program, the invention can identify and display related news stories. The invention can also enable the user to control the display of the news programs by, for example, speeding up the display, causing a summary of one or more news stories to be displayed, or pausing the display of the news stories. Additionally, the invention can indicate to the user which news story is currently being viewed, as well as which news stories have previously been viewed.

129 Claims, 6 Drawing Sheets



US 6,263,507 B1

Page 2

U.S. PATENT DOCUMENTS

5,524,051	6/1996	Ryan	380/9
5,537,151 *	7/1996	Orr et al.	348/564
5,590,195	12/1996	Ryan	380/9
5,614,940 *	3/1997	Cobbley et al.	348/7
5,689,648	11/1997	Diaz et al.	455/3.1
5,703,795 *	12/1997	Mankovitz	348/473
5,751,806	5/1998	Ryan	380/9
5,774,664 *	6/1998	Hidary et al.	348/12
5,892,536	4/1999	Logan et al.	348/13
6,020,883	2/2000	Herz et al.	345/327
6,025,837	2/2000	Mathews, III et al.	345/327
6,061,056	5/2000	Menard et al.	345/327

OTHER PUBLICATIONS

Elliot, E., "Watch, Grab, Arrange, See: Thinking with Motion Images via Streams and Collages", Masters thesis, School of Architecture and Planning, Massachusetts Institute of Technology, Feb. 1993, pp. 3, 5, 7, 9-11, 13-35, 37-49, 51-61, 63-85, 87-99, 101, 103-105.

CNN At Work White Paper, 1994.

Sahon, G. et al., "Improving Retrieval Performance by Relevance Feedback", Journal of the American Society for Information Science, vol. 41, No. 4, Jun. 1990, pp. 288-297.

Shibata M., "A Description Model of Video Content and Its Application for Video Structuring", Systems and Computers in Japan, vol. 27, No. 7, Jun. 1996, pp. 70-83.

Yeung M. M. et al., "Efficient Matching and Clustering of Video Shots", IEEE '95 (ICIP), vol. 1, Oct. 1995, pp. 338-341.

Wactlar H. D. et al., "Intelligent Access to Digital Video: Informedia Project", Computer, vol. 29, No. 5, May 1996, pp. 46-52.

Sakauchi M. et al., "Multimedia Database Systems for the Contents Mediator", IEICE Trans. Inf. and Syst., vol. E79-D, No. 6, Jun. 1996, pp. 641-646.

Shahraray, B. et al., "Automatic generation of pictorial transcripts of video programs", SPIE, vol. 2417, Jan. 1995, pp. 512-518.

Lindblad C.J. et al., "ViewStation Applications: Implications for Network Traffic", IEEE Journal on Selected Areas in Communications, vol. 13, No. 5, Jun. 1995, pp. 768-777.

C. Horner, "NewsTime: A Graphical User Interface to Audio News," Masters thesis, School of Architecture and Planning, Massachusetts Institute of Technology, Jun. 1993, pp. 1-84. U.S. application No. 08/528,891, Ahmad, filed Sep. 15, 1995.

U.S. application No. 08/399,482, Covell et al., filed Mar. 17, 1995.

A. Hauptmann et al., "Text, Speech, and Vision for Video Segmentation: The Informedia™ Project," AAAI Fall Symposium, Computational Models for Integrating Language and Vision, Nov. 10-12, 1995.

D. Tennenhouse et al., "The ViewStation: a software-intensive approach to media processing and distribution," Proceedings of the 17th Annual Conference on Research and Development in Information Retrieval, Jul. 3-6, 1994, pp. 104-113.

H. Zhang et al., "Automatic Parsing of News Video," IEEE Conference on Multimedia Computing and Systems, 1994, pp. 45-54.

E. Scheirer et al., "Construction and Evaluation of a Robust Multifeature Speech/Music Discriminator," Proc. ICASSP, Apr. 21-24, 1997, pp. 1-4.

C. Buckley et al., "The Effect of Adding Relevance Information in a Relevance Feedback Environment," Proceedings of the 17th Annual Conference on Research and Development in Information Retrieval, Jul. 3-6, 1994, pp. 292-300.

S. Roucos et al., "High Quality Time-Scale Modification for Speech," Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, IEEE, 1985, pp. 493-496.

D. Reynolds, "A Gaussian Mixture Modeling Approach To Text-Independent Speaker Identification," Ph.D. thesis, Dept. of Electrical Engineering, Georgia Institute of Technology, 1992, pp. 1-154.

* cited by examiner




FIG. 1

U.S. Patent

Jul. 17, 2001

Sheet 2 of 6

US 6,263,507 B1




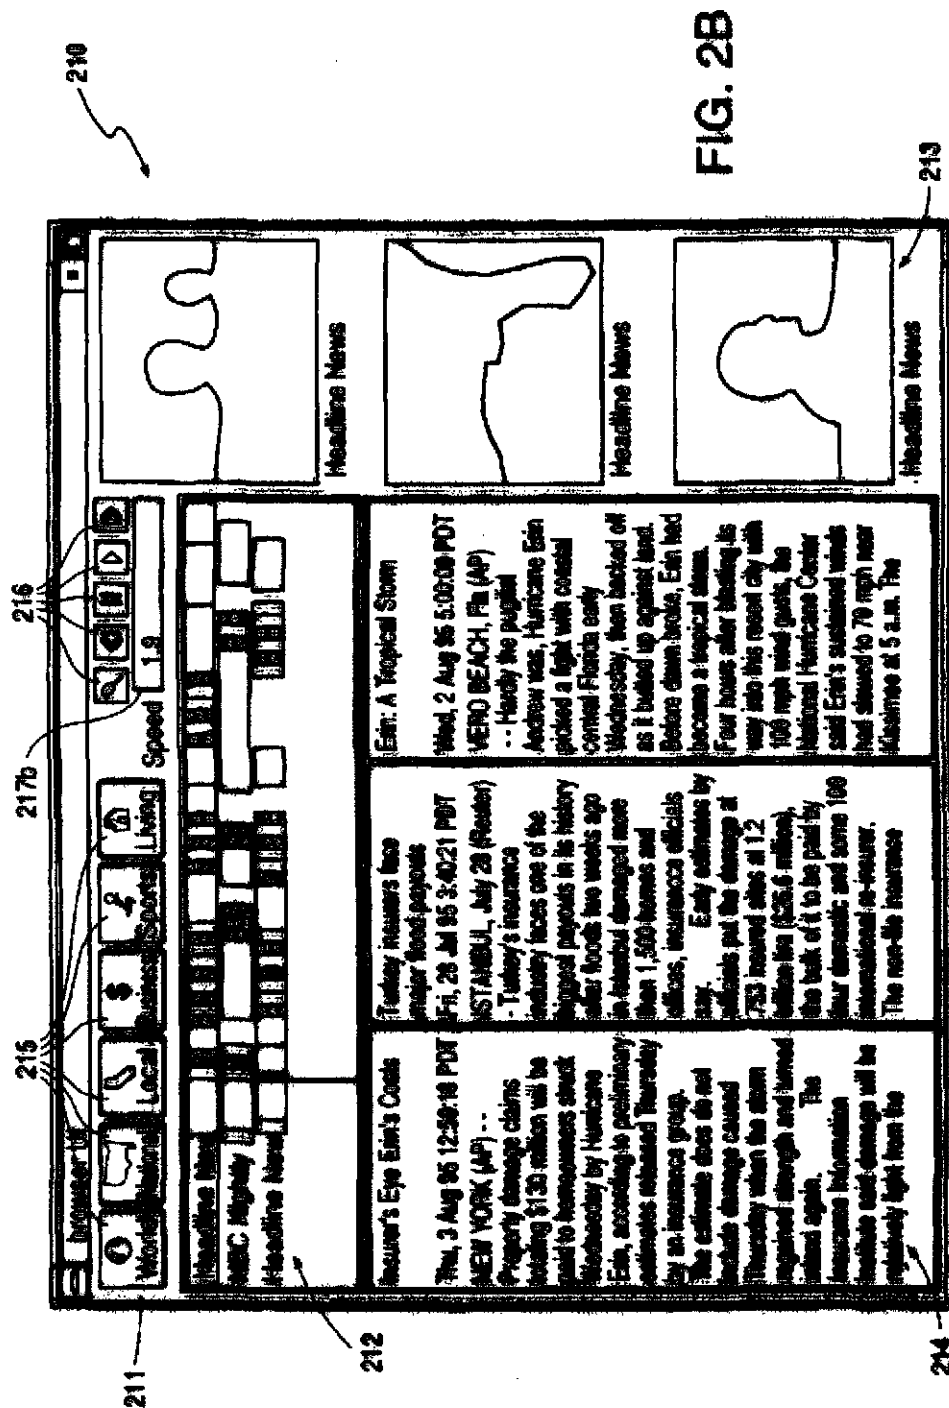
FIG. 2A

U.S. Patent

Jul. 17, 2001

Sheet 3 of 6

US 6,263,507 B1



U.S. Patent

Jul. 17, 2001

Sheet 4 of 6

US 6,263,907 B1

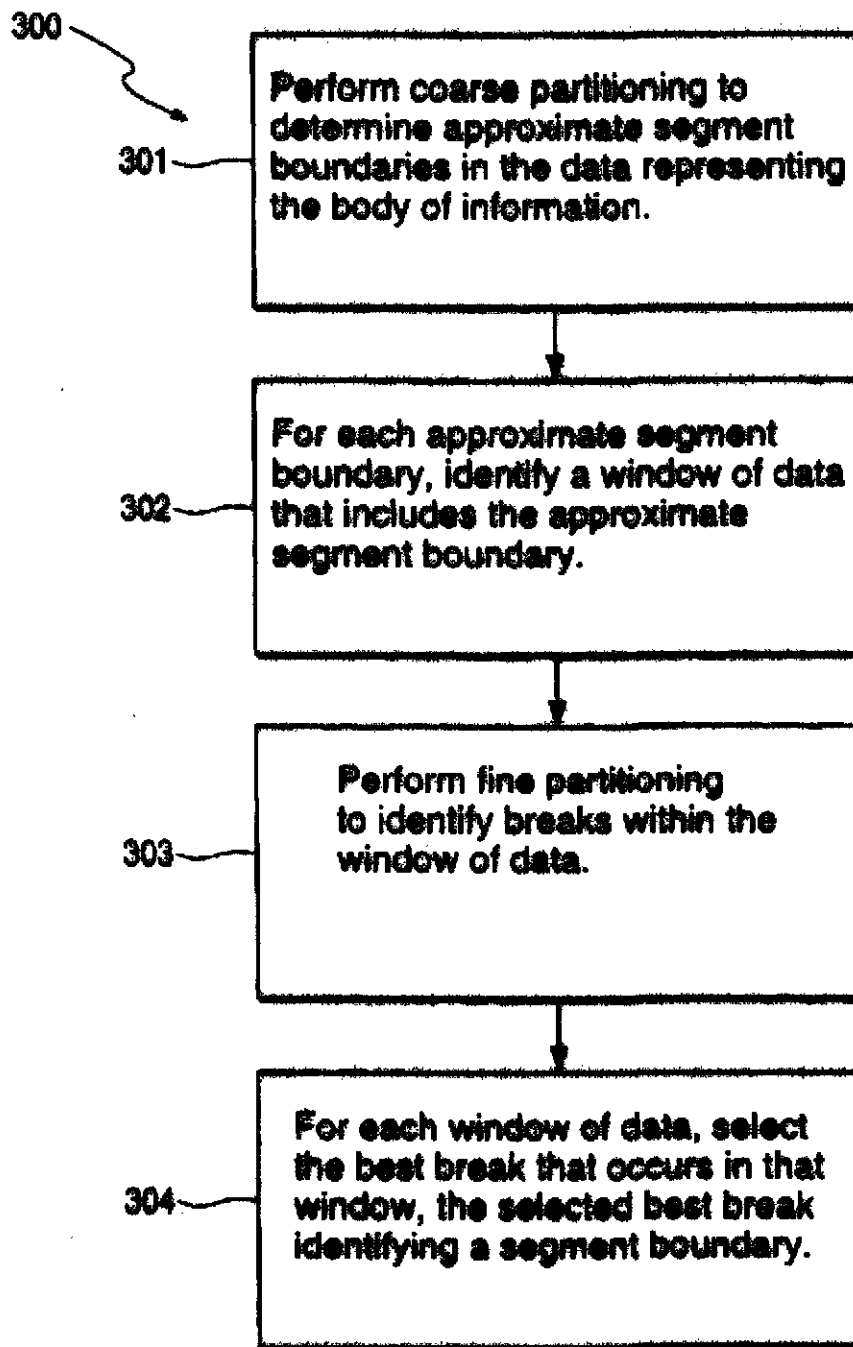


FIG. 3

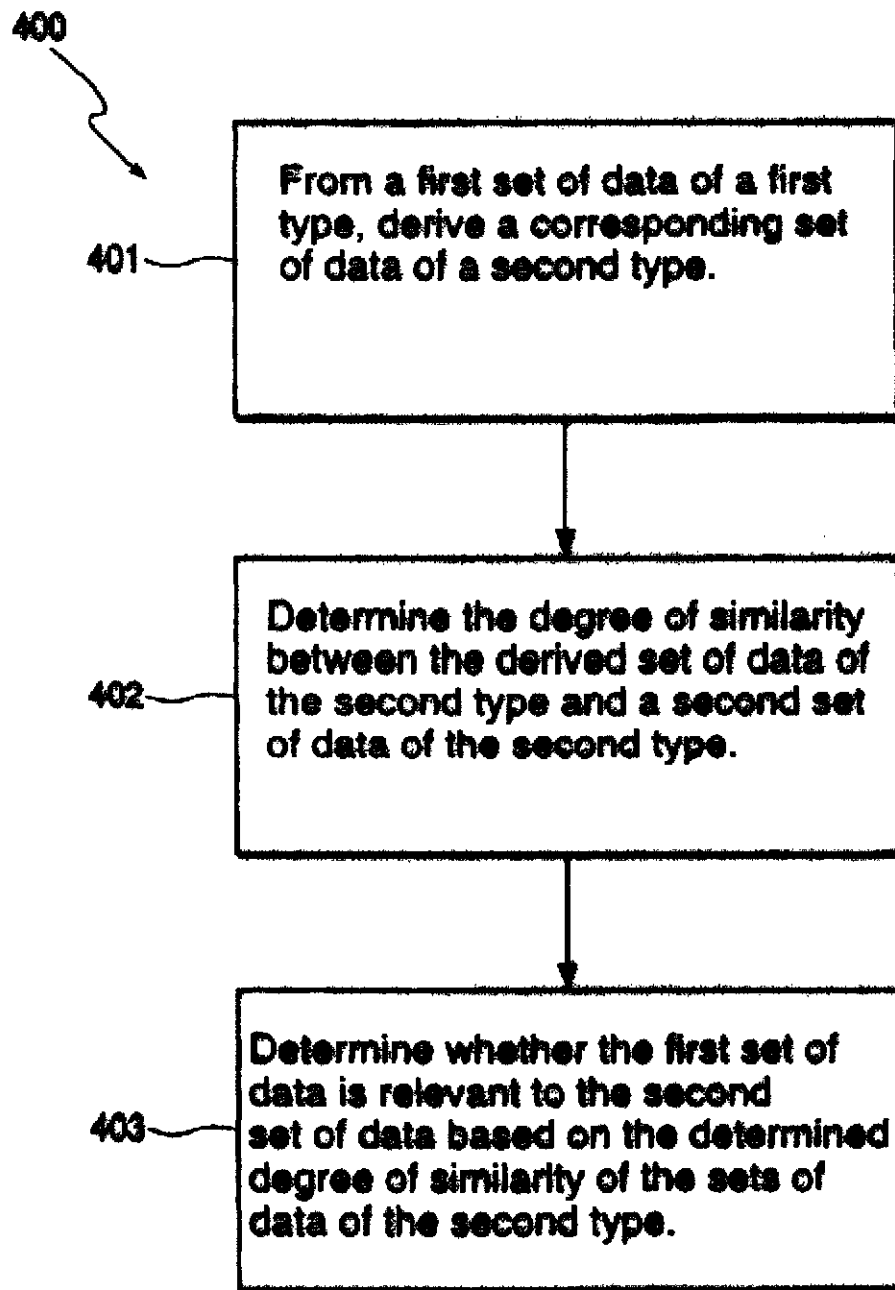


FIG. 4

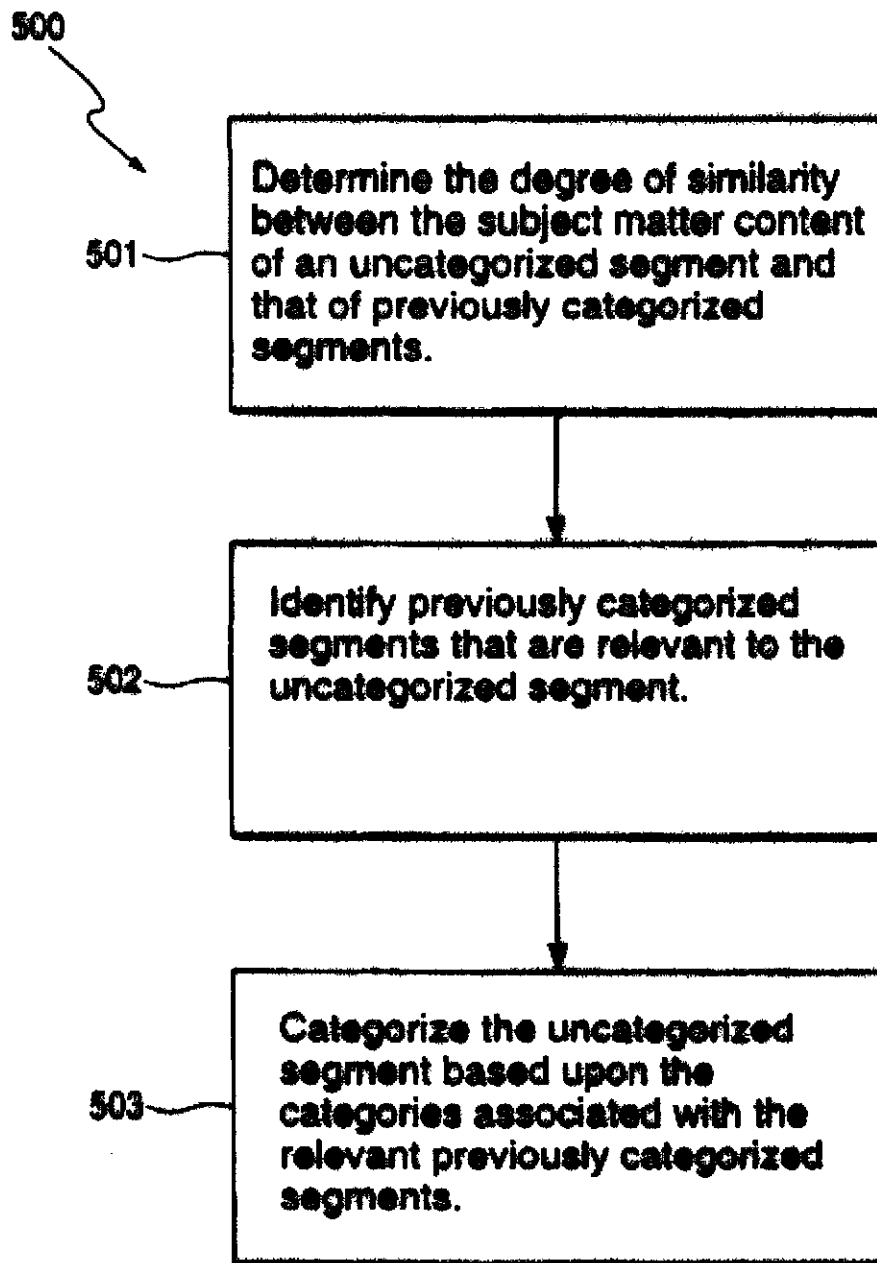


FIG. 5

US 6,263,507 B1

1

BROWSER FOR USE IN NAVIGATING A BODY OF INFORMATION, WITH PARTICULAR APPLICATION TO BROWSING INFORMATION REPRESENTED BY AUDIOVISUAL DATA

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to systems and methods that enable observation of a body of information and, in particular, a body of information that can be represented, at least in part, by audiovisual data. Most particularly, the invention relates to systems and methods for accessing and reviewing a body of information represented by one or more sets of audiovisual data that can be used to generate an audiovisual display and one or more related sets of text data that can be used to generate a text display.

2. Related Art

The increasing complexity of the modern world, and the concomitant explosion in the amount of information available to describe that world, has placed competing demands on people. There is more subject matter that people find necessary or desirable to master or, at least, be familiar with. At the same time, there is less time to spend delving into any particular subject. Too, there is a much larger universe of information from which the desired information must be extracted. Trying to get just an overview of a large body of information can be overwhelming, and attempting to find specific material within the body of information can be like searching for a needle in a haystack.

Thus, there is a continuing and growing need for methods and systems for enabling bodies of information to be accessed and reviewed in a useful manner, e.g., a manner that allows the scope and content of available information to be quickly ascertained and that enables quick access to information of particular interest. In particular, there is a need for systems and methods of organizing, categorizing and relating the various segments of a large body of information to facilitate the access and review of the body of information. For example, while some previous systems for enabling observation of a large body of information enable identification of one or more segments of information that are related to a specified segment of information, these systems do not automatically display such related segments of information. Moreover, the previous systems either require that related segments have previously been determined or, at least, that the segments have been categorized according to subject matter content so that whether two segments are related can readily be determined. Further, previous systems have not enabled determination of relatedness between segments of information represented by different types of data, e.g., such systems cannot determine whether a segment represented by audiovisual data is related to a segment represented by text data.

There is also a need for systems and methods for enabling observation of a body of information that are user-friendly, e.g., that can be used with little training, that are convenient to use, that enable information to be quickly and easily accessed, and that present the information in an accessible format via a high quality display medium. It would also be desirable for such systems and methods to be adapted for use with bodies of information represented by different types of data (i.e., audio data, video data, text data or some combination of the three). It would further be desirable for such systems and methods to be adapted for use with bodies of information represented by data acquired from a wide vari-

2

ety of media (e.g., print media such as newspapers or magazines, television and radio broadcasts, online computer information services and pre-recorded audiovisual programs, to name a few). Previous systems and methods for accessing and reviewing a body of information are deficient in one or more of these respects.

For example, many previous systems are computer-based. Typically, the display device of these systems (e.g., conventional computer display monitor) does not provide a high quality display of time-varying audiovisual information (such as produced by a television, for example). On the other hand, display devices that do display such information well (e.g., televisions), typically do not provide a high quality display of text information (such as produced by a computer display monitor). A system that can provide a high quality display of both types of information is needed.

Additionally, previous systems for reviewing a body of information are not as flexible or convenient to use as is desirable. For example, in many such systems (e.g., computers), the mechanism for controlling the operation of the system is physically coupled to the display device of the system. Therefore, the system can not be operated remotely, thus constraining the user's freedom of movement while operating the system. Additionally, even in those systems where remote operation is possible (e.g., remotely controlled televisions), the remote control device often does not have a user interface that is as readily accessible as desired (as many consumer electronics users can testify, the keypads of many remote control devices are an impenetrable array of cryptic control keys, often requiring non-intuitive key combinations to effect particular control instructions) or the remote control device does not contain a rich set of control features. Moreover, the remote control devices used with previous systems do not have the capability of themselves displaying a part of the body of information.

Further, previous systems often do not enable real-time acquisition and review of some or all of the body of information. For example, many computer-based systems acquire and store data representing a body of information. The stored data can then be accessed to enable display of segments of the body of information. However, insofar as previous systems for observing a body of information allow real-time acquisition and review of the body of information, these systems generally do not analyze the data to enable the data to be organized, categorized and related so that, for example, segments of the body of information can be related to other segments for which data is acquired in the future or for which data has previously been acquired. Moreover, such systems do not enable the real-time display of some or all of a body of information while also displaying related information in response to the real-time display.

Thus, there is a need for improved systems and methods for enabling observation of a body of information and, in particular, such systems and methods that address the above-identified inadequacies in previous systems and methods for enabling observation of a body of information.

SUMMARY OF THE INVENTION

The invention enables a body of information to be displayed by electronic devices (e.g., a television, a computer display monitor) in a manner that allows the body of information to be reviewed quickly and in a flexible manner. Typically, the body of information will be represented by a set of audio data, video data, text data or some combination of the three. In a particular embodiment, the invention enables generation of an audiovisual display of one or more

US 6,263,507 B1

3

segments of information, as well as a display (a text display, an audio display, a video display, or an audiovisual display), for each of the segments, of one or more related segments of information. In a particular application of the invention, referred to herein as a "news browser", the invention enables acquisition, and subsequent review, of news stories obtained over a specified period of time from a specified group of news sources. For example, as a news browser, the invention can be used to review news stories acquired during one day from several television news programs (e.g., CNN Headline News, NBC Nightly News), as well as from text news sources (e.g., news wire services, traditional print media such as newspapers and magazines, and online news services such as Clarinet™).

The invention enables some or all of a body of information to be skimmed quickly, enabling a quick overview of the content of the body of information to be obtained. The invention also enables quick identification of information that pertains to a particular subject. The invention further enables quick movement from one segment of a body of information to another, so that observation of particular information of interest can be accomplished quickly. In a news browser according to the invention, for example, each of a set of television news programs can be skimmed to quickly ascertain the subject matter content of the news stories contained therein. Additionally, a particular category (e.g., subject matter category) can be specified and news stories having content that fits within the specified subject matter category can be immediately identified and either displayed or identified as pertinent to the subject matter category and available for display. Further, a user of the news browser can move arbitrarily among news stories within the same or different news programs.

The invention also enables automatic identification of information that is related to information that is being displayed, so that the related information can be observed, thereby enabling information about a particular subject to be examined in depth. In particular, the invention enables such identification of related segments to be made between segments of different types (e.g., a segment represented by audiovisual data can be compared to a segment represented by text data to enable a determination of whether the segments are related). A portion or a representation of the related information can be displayed in response to (e.g., simultaneous with) the original information display. For instance, in a news browser according to the invention, one or more text news stories (e.g., news stories that are obtained from traditional print media or from electronic publications) that are related (i.e., which cover the same or similar subject matter) to a television news story being displayed can be automatically identified and a portion of the related text news story or stories displayed so that the story or stories can be reviewed for additional information regarding the subject matter of the television news story. Additionally, in a news browser according to the invention, one or more other television news stories that are related to a television news story being displayed can be automatically identified and a single representative video frame displayed for each such news story.

Additionally, the invention enables automatic categorization of uncategorized segments of the body of information based upon comparison to other segments of the body of information that have been categorized. In particular, the subject matter category of a segment of information can be determined by comparing the segment to one or more previously categorized segments and categorizing the segment in accordance with the subject matter categorization of

4

one or more previously categorized segments that are determined to be relevant to the uncategorized segment. In a news browser according to the invention, for example, this can be used to categorize the news stories of a television news program based upon the categorization of text news stories that are found to be relevant to the television news stories.

The invention can be implemented in a system that is convenient to use, that presents the body of information in a readily accessible way, and that presents the information via one or more display devices that are tailored for use with the particular type of data that is used to generate the display. For example, a system according to the invention can include a control device that enables remote, untethered control of a primary display device of the system. The remote control device can also be implemented so that some or all of the body of information can also be displayed on the remote control device. The system can include, for example, a television for display of audiovisual information and a computer display monitor for display of text information.

Additionally, a control device of a system according to the invention can be implemented with a graphical user interface that facilitates user interaction with the system. For example, such an interface can include a region that provides an indication of a user's past progression through, and present location within, the body of information. In a news browser according to the invention, for example, a program map is displayed that facilitates navigation through the news programs that can be selected for display.

The invention also enables real-time acquisition and review of some or all of the body of information. The invention enables on-the-fly analysis of data as the data is acquired, so that the data can be organized, categorized and related to other data. The invention also enables the real-time display of some or all of a body of information while also displaying related information in response to the real-time display. For example, in a news browser according to the invention, television news programs can be acquired and displayed as they occur. Related news stories, either from previously acquired television news programs or text news sources can be displayed as each television news story is displayed in real time.

The invention also enables control of the manner in which the information is displayed (e.g., the apparent display rate of the display can be controlled, the display can be paused, a summary of a portion of the body of information can be displayed). For example, in a news browser according to the invention, the user can cause a summary of one or more television news stories to be displayed (rather than the entire news story or stories), the user can speed up (or slow down) the display of a television news story, and the user can pause and resume the display of a television news story such that the display resumes at an accelerated rate until the display of the news story "catches up" to where the display would have been without the pause (a useful feature when the television news story is being acquired and displayed in real time).

In one aspect of the invention, a system enables acquisition and review of a body of information that includes a multiplicity of segments that each represent a defined set of information (frequently, a contiguous related set of information) in the body of information. The system includes: i) a mechanism for acquiring data representing the body of information; ii) a mechanism for storing the data; iii) a first display mechanism for generating a display of a first segment of the body of information from data that is part of the stored data; iv) a mechanism for comparing the data representing a segment of the body of information to the data

US 6,263,507 B1

5

representing a different segment of the body of information to determine whether, according to one or more predetermined criteria, the compared segments are related; and v) a second display mechanism for generating a display of a portion of, or a representation of, a second segment of the body of information from data that is part of the stored data. (A method according to the invention, and a computer readable medium encoded with one or more computer programs according to the invention, both enable similar capability.) The second display mechanism displays a portion or representation of the second segment in response to the display by the first display mechanism of a first segment to which the second segment is related. The second display mechanism can display a portion or representation of the second segment substantially coextensive in time with the display of the related first segment by the first display mechanism. The system can further include a mechanism for identifying the subject matter content of a segment of the body of information, so that the mechanism for comparing can determine the similarity of the subject matter content of a segment to the subject matter content of a different segment (using, for example, relevance feedback) and use that result to determine the relatedness of the compared segments. The system can also include a mechanism for identifying an instruction from a user to begin displaying at least some of the body of information, the first display mechanism beginning display of a segment in response to the user instruction. When a portion or representation of a second segment is being displayed, the system can enable such a second segment to be selected for display by the first display mechanism. Often, the segments displayed by the first display mechanism are represented by audiovisual data (and, in particular, audiovisual data that can be used to generate an audiovisual display that can vary with time), such as, for example, data produced from television or radio broadcast signals. The segments displayed by the second display mechanism can be represented by audiovisual data (e.g., a single representative video image, or "keyframe") or by text data (e.g., text excerpts), such as, for example, data from computer-readable data files acquired over a computer network from an information providing site that is part of that network. In particular applications for which use of the invention is contemplated, the first display mechanism can be an analog display device (such as a television) and the second display means can be a digital display device (such as a computer display monitor). The system can advantageously be implemented so that the various devices are interconnected to a conventional computer bus that enables the devices to communicate with each other such that the devices do not require wire communication over network communication lines to communicate with each other (the devices are "untethered").

In another aspect of the invention, a system for reviewing a body of audiovisual information that can vary with time (e.g., the content from one or more news broadcasts) includes: i) a mechanism for displaying the audiovisual information; and ii) a mechanism for controlling operation of the system, the mechanism for controlling being physically separate from the mechanism for displaying and including a graphical user interface for enabling specification of control instructions. The mechanism can advantageously be made portable. Further, the system can advantageously include a mechanism for 2-way wireless communication between the mechanism for displaying and the mechanism for controlling. The graphical user interface can include one or more of the following: i) a playback control region for enabling specification of control instruc-

6

tions that control the manner in which the audiovisual information is displayed on the means for displaying; ii) a map region for providing a description of the subject matter content of the audiovisual information and for enabling specification of control instructions that enable navigation within the audiovisual information; iii) a related information region for displaying a portion of, or a representation of, a segment that is related to a segment being displayed by the mechanism for displaying; and iv) a secondary information display region for displaying a secondary information segment that is related to a segment of the audiovisual information that is being displayed by the mechanism for displaying. In particular, the playback control region can include one or more of the following: i) an interface that enables selection of one of a plurality of subject matter categories, all of the segments of the audiovisual information corresponding to a particular subject matter category being displayed in response to the selection of that subject matter category; ii) an interface that enables variation of the apparent display rate at which the audiovisual information is displayed; iii) an interface that enables specification of the display of a summary of a segment of the audiovisual information; iv) an interface that enables the display to be paused, then resumed at an accelerated rate that continues until the display of the audiovisual information coincides with the display that would have appeared had the display not been paused; v) an interface that enables termination of the current segment display and beginning of a new segment display; and vi) an interface that enables repetition of the current segment display. The map region can further identify a segment of the audiovisual information that is currently being displayed and/or identify each segment of the audiovisual information that has previously been displayed.

In still another aspect of the invention, a system enables review of a body of information, the body of information including a first portion that is represented by audiovisual data that can vary with time and a second portion that is represented by text data. The system includes a first display device for displaying the first portion of information and a second display device for displaying the second portion of information. The first display device is particularly adapted for generation of a display from time-varying audiovisual data, while the second display device is particularly adapted for generation of a display from text data. The first display device can be, for example, an analog display device such as a television. The second display device can be, for example, a digital display device such as a computer display monitor. The two devices can interact with each other so that related information can be displayed at the same time on the two devices, in the same manner as that described above.

In another aspect of the invention, a method categorizes according to subject matter a segment of a body of information (that includes a plurality of segments), the segment not previously having been categorized according to subject matter, based upon the subject matter category or categories associated with one or more previously categorized segments of the body of information. The uncategorized segment can have been acquired from a first data source (that supplies, for example, television or radio broadcast signals) and the previously categorized segment or segments can have been acquired from a second data source (that supplies, for example, computer-readable data files) that is different than the first data source. The method includes the steps of: i) determining the degree of similarity between the subject matter content of the uncategorized segment and the subject matter content of each of the previously categorized segments; ii) identifying one or more of the previously catego-

US 6,263,507 B1

7
 sized segments as relevant to the uncategorized segment based upon the determined degrees of similarity of subject matter content between the uncategorized segment and the previously categorized segments; and iii) selecting one or more subject matter categories with which to identify the uncategorized segment based upon the subject matter category or categories used to identify the relevant previously categorized segment or segments. (A computer readable medium encoded with one or more computer programs according to the invention enables similar capability.) The step of determining the degree of similarity can be accomplished using a relevance feedback method. The step of identifying one or more of the previously categorized segments as relevant to the uncategorized segment can include the steps of: i) identifying a multiplicity of the previously categorized segments that are the most similar to the uncategorized segment; ii) determining the degree of similarity between each of the multiplicity of previously categorized segments and each other of the plurality of previously categorized segments; iii) for each pair of previously categorized segments of the multiplicity of previously categorized segments having greater than a predefined degree of similarity, eliminating one of the pair of previously categorized segments from the multiplicity of previously categorized segments, wherein the previously categorized segment or segments remaining after the step of eliminating are similar and distinct previously categorized segments; and iv) identifying one or more of the similar and distinct previously categorized segments as relevant previously categorized segments.

In another aspect of the invention, a method determines whether a first set of information represented by a set of data of a first type (e.g., text data) is relevant to a second set of information (that is different than the first set of information) represented by a set of data of a second type (e.g., audiovisual data). The method includes the steps of: i) deriving a set of data of the second type from the set of data of the first type, the derived set of data of the second type also being representative of the first set of information; ii) determining the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information; and iii) determining whether the first set of information is relevant to the second set of information based upon the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information. (A computer readable medium encoded with one or more computer programs according to the invention enables similar capability.) The step of determining the degree of similarity can be accomplished using a relevance feedback method. Still further in accordance with this aspect of the invention, a method can determine which, if any, of a multiplicity of sets of information represented by an associated set of data of a first type (each of the multiplicity of sets of information being different from other of the multiplicity of sets of information) are relevant to the second set of information represented by the set of data of the second type. This method includes the steps of, in addition to those discussed above: i) determining the degree of similarity between each set of data of the first type representing one of the multiplicity of sets of information and the derived set of data of the first type representing the second set of information; ii) identifying which, if any, of the sets of data of the first type representing one of the multiplicity of sets of information have greater than a predefined degree of similarity to the derived set of data of the first type

8
 representing the second set of information, the sets of data of the first type so identified being termed similar sets of data of the first type; iii) determining the degree of similarity between each similar set of data of the first type and each other similar set of data of the first type; iv) for each pair of similar sets of data of the first type having greater than a predefined degree of similarity, eliminating one of the pair of similar sets of data of the first type from the set of similar sets of data of the first type, wherein the set or sets of similar data of the first type remaining after the step of eliminating are similar and distinct sets of data of the first type; and v) identifying the set or sets of information corresponding to one or more of the similar and distinct sets of data of the first type as relevant to the second set of information.

In still another aspect of the invention, a method enables the identification of the boundaries of segments in a body of information that is represented by a set of text data and at least one of a set of audio data or a set of video data, each segment representing a contiguous related set of information in the body of information. (A computer readable medium encoded with one or more computer programs according to the invention enables similar capability.) The segment boundaries are identified by first performing a coarse partitioning method to approximately locate the segment boundaries, then performing a fine partitioning method to more precisely locate the segment boundaries. In the coarse partitioning method, time-stamped markers in the set of text data are identified and used to determine approximate segment boundaries within the body of information. For each time of occurrence of an approximate segment boundary in the text data, a range of time is specified that includes the time of occurrence. Subsets of audio data or subsets of video data that occur during the specified ranges of time are extracted from the complete set of audio data or the complete set of video data. The fine partitioning method is then performed to identify one or more breaks in each of the subsets of audio data or each of the subsets of video data. The best break that occurs in each subset of audio data or each subset of video data is selected, and the time of occurrence of the best break in each subset is designated as a boundary of a segment in the body of information. The fine partitioning can be performed using any appropriate method. For example, when segment boundaries are being determined in video data, scene break identification can be used to implement the fine partitioning. When segment boundaries are being determined in audio data, the fine partitioning can be implemented by, for example, pause recognition, voice recognition, word recognition or music recognition. Once segment boundaries have been determined in the audio data or the video data, a synchronization of the audio data and the video data can be used to determine the boundaries of the segment in the other of the audio data or video data.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating a system according to the invention for acquiring and reviewing a body of information.

FIG. 2A is a diagrammatic representation of a graphical user interface according to the invention that can be used to enable control of the operation of a system according to the invention, display information regarding operation of the system of the invention and display information acquired by the system of the invention.

FIG. 2B is a view of an illustrative graphical user interface in accordance with the diagrammatic representation of FIG. 2A.

US 6,263,507 B1

9

FIG. 3 is a flow chart of a method in accordance with the invention for identifying the boundaries of segments in a body of information.

FIG. 4 is a flow chart of a method in accordance with the invention for determining whether a first set of information represented by data of a first type is relevant to a second set of information represented by data of a second type.

FIG. 5 is a flow chart of a method in accordance with the invention for categorizing according to subject matter an uncategorized segment of a body of information based on the categorization of other previously categorized segments of the body of information.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

I. Overview

Generally, the invention enables the acquisition of a body of information and review of the content of the body of information. In particular, the invention includes various features that facilitate and enhance review of the body of information. The invention enables the body of information to be quickly reviewed to obtain an overview of the content of the body of information or some portion of the body of information. The invention also allows flexibility in the manner in which the body of information is reviewed. For example, the invention enables a user to move quickly from one segment of a body of information to another, enabling the user to rapidly begin observing particular information of interest. Further, the invention enables a user to quickly locate information within the body of information that pertains to a particular subject in which the user has an interest. The invention also enables a user to, when observing particular information, quickly find and review other information that is related to the information that the user is observing. Additionally, the invention enables the user to control the manner in which the information is displayed (e.g., the apparent display rate of the display can be controlled, the display can be paused, a summary of a portion of the body of information can be displayed). The invention also provides the user with an indication of the user's past progression through, and present location within, the body of information, such indications aiding the user in selecting further segments (described below) of the body of information for review.

The body of information can be represented by one or more sets of audio data, one or more sets of video data, one or more sets of text data or some combination of the three. Herein, "audio data" refers to data used to generate an audio display, "video data" refers to data used to generate a video display substantially including images other than text images, "text data" refers to data used to generate a video (or audio, though typically video) display of text images, and "audiovisual data" refers to data that includes audio and/or video data, and may include text data. In a particular embodiment, the invention enables the acquisition and review of one or more sets of information represented by audiovisual data, as well as related sets of information represented by text data.

For example, in a particular application of the invention, the content of one or more audiovisual news programs is acquired from a first set of one or more information sources and news stories (or "articles") from text news sources are acquired from a second set of one or more information sources. The first set of information sources could be, for example, CNN Headline News or network (e.g., ABC, NBC,

10

CBS) news programs. The second set of information sources could be, for example, on-line news services such as Clarinet™ or news wire services such as AP or UPI. It is contemplated that this application of the invention can be particularly useful as a means of enhancing the viewing of conventional television news programs. For example, in this application, the invention can enable the user to access the news stories of audiovisual news programs in a random manner so that the user can move quickly from one news program to another, or from one news story in a news program to another news story in the same or another news program. The invention can also enable the user to quickly locate news stories pertaining to a particular subject. Additionally, when the user is observing a particular news story in an audiovisual news program, the invention can identify and display a related text news story or stories. The invention can also enable the user to control the display of the audiovisual news programs by, for example, speeding up the display, causing a summary of one or more news stories to be displayed, or pausing the display of the news stories, thereby enabling the user to quickly ascertain the content of one or more news stories or entire news programs. Additionally, the invention can indicate to the user which audiovisual news program is currently being viewed (and, further, which news story within the news program is being viewed), as well as which news stories and/or news programs have previously been viewed.

II. System Configuration

FIG. 1 is a block diagram illustrating a system 100 according to the invention for acquiring and reviewing a body of information. A user 109 interacts with a control device 101 to cause information to be displayed on a primary display device 102. The control device 101 includes an appropriate user interface (e.g., a graphical user interface, as discussed in more detail below) that allows the user 109 to specify control instructions for effecting control of the system 100. Communication between the control device 101 and the primary display device 102 is mediated by a system controller 103. The system controller 103 causes primary information to be acquired from a primary information source 107 via a primary information data acquisition device 105. Herein, "primary information" is any information the display of which the user can directly control. The system controller 103 also causes secondary information (which is typically related to the primary information) to be acquired from a secondary information source 108 via a secondary information data acquisition device 106. Herein, "secondary information" is any information other than primary information that is acquired by a system according to the invention and that can be displayed by the system and/or used by the system to manipulate or categorize (as described in more detail below) the primary information. A data storage device 104 stores the acquired primary and secondary information. The primary information is displayed on the primary display device 102. The secondary information can be displayed (e.g., by the control device 101 or by the primary display device 102 in addition to the primary information) or not (i.e., the secondary information may be used only for categorizing and/or manipulation of the primary information). Illustratively, the primary information can be videotape (or other audiovisual data representation) of an audiovisual news program or programs and the secondary information can be the text of news stories from text news sources.

The control device 101, the primary display device 102, the system controller 103 and the data storage device 104

US 6,263,507 B1

11

can be embodied in one or more devices that can be interconnected to a conventional computer bus that enables the devices to communicate with each other. In particular, the devices 101, 102, 103 and 104 can be integrated into a system in which the devices do not require wire communication over network communication lines to communicate with each other (one or more of the devices 101, 102, 103 and 104 is "untethered" with respect to one or more of the other devices 101, 102, 103 and 104). Thus, once the primary and secondary information have been acquired by the system 100, the primary and secondary information can be accessed and displayed at a relatively fast speed, thus providing quick response to control instructions from the user and enabling generation of displays with acceptable fidelity. In contrast, a networked system in which the devices must communicate with each other over a network via wire communication lines—in particular, a system in which the control device and display device or devices must communicate over such wire communication lines with the data storage device on which the information is stored—may not produce acceptable performance. In the networked system, the operation of the system is limited by the communications bandwidth and latency of the network communications medium. For example, the bandwidth of the network communications medium may not be adequate to enable transfer of data from the data storage device 104 to the primary display device 102 quickly enough to enable a display with acceptable fidelity to be generated by the primary display device 102. Or, the response to a control instruction from the control device 101 may be undesirably slow because of inadequate speed of the network communications medium.

The primary information data acquisition device 105 and secondary information data acquisition device 106 can be implemented by any appropriate such devices. Where the primary information source 107 is comprised of television news broadcasts, for example, the primary information data acquisition device 105 can be a conventional television tuner and video capture device that acquires the data representing the primary information via conventional cable connections, satellite dish or television antenna. Where the secondary information is comprised of online text sources (i.e., text sources available over a computer network such as the Internet), for example, the secondary information data acquisition device 106 can be a conventional modem or other communications adapter, as known by those skilled in the art of data communications, that enables acquisition of data representing the secondary information via one or more conventional communication lines, such as telephone lines, ISDN lines or Ethernet connections. (It is also possible that the primary information can be acquired from online sources, such as via the Internet or other computer network.)

The primary information data acquisition device 105 and the secondary information data acquisition device 106 can communicate with the system controller 103 in any appropriate manner. As described below, the system controller 103 can be implemented as part of a digital computer. Where this is the case, the communication between the system controller 103 and the devices 105 and 106 is preferably implemented to enable computer control of the devices 105 and 106. When the device 105 or 106 is used to acquire information over a computer network, the device 105 or 106 will be a device, such as a computer modem, for which such communication to the system controller 103 can be implemented using well-known methods and apparatus. For other types of devices, such communication must be implemented in another manner. For example, when the device 105 is a television tuner, communication between the system con-

12

troller 103 and the device 105 can be implemented using a VISCA (Video System Control Architecture) connection.

As will be apparent from the description below, the processing of the data representing the primary and secondary information generally requires that the data be in digital form. Text data acquired from online text sources, for example, is acquired in digital form and so can be used directly in such processing. Analog television signals, however, must be digitized before being used in digital processing. This can be accomplished using conventional A/D conversion methods and apparatus. Further, it is desirable to compress the data to increase the amount of data (i.e., primary and secondary information) that can be stored on the data storage device 104. For example, the television data can be compressed according to the MPEG, JPEG or MJPEG video compression standards, as known by those skilled in the art of audio and video data compression. The text data can also be compressed, using conventional text file compression programs, such as PKZIP, though, typically, such compression provides a relatively small benefit because the amount of text data is small compared to the amount of audio and video data, and the amount of data required to represent the categorization information (described below). Finally, it may be desirable or necessary to transform digital data into an analog waveform again (e.g., convert digital video data into analog video data for display by a television). This can be accomplished using conventional D/A conversion methods and apparatus.

In the embodiment of the invention shown in FIG. 1, the system 100 according to the invention makes use of two devices for display and control: a primary display device 102 for displaying the primary information and a control device 101 for controlling the operation of the primary display device 102. Preferably, the control device 101 is physically separate from the primary display device 102 and portable so that the user has flexibility in selecting a position relative to the primary display device 102 during use of the system 100. For example, such an embodiment could allow a user to use the invention while sitting in a chair or on a couch, reclining in bed, or sitting at a table or desk. Additionally, when the secondary information is textual (e.g., the text of news stories) and the control device 101 is used to display such secondary information, the portability of the control device 101 attendant such an embodiment increases the likelihood that the text is displayed on a device that can be held in close proximity to the user, thereby improving the ability of the user to view the text. Further, as discussed in greater detail below, the control device 101 preferably has sophisticated user interface capabilities.

As previously mentioned, a system according to the invention (including the system 100) can be implemented so that the primary display device 102 displays the primary information while a separate device (e.g., the control device 101) displays the secondary information. Further, as can be appreciated from the description herein, the invention can advantageously be used in situations in which the primary information is audiovisual information (and, in particular, audiovisual information that can vary with time, such as the content of a television program) and the secondary information is text information (some or all of which is, typically, likely to be related to the audiovisual information). In such an implementation of the invention, the use of two different devices for display allows the optimization of the display devices for the particular type of information to be displayed. (A system according to the invention can, in general, have any number of displays, as necessary or advantageous.) Thus, where the primary information is audiovisual

US 6,263,507 B1

13

information, the primary display device 102 is preferably a device that enables high quality audio and video images (in particular, time-varying audio and video images) to be produced, such as a television. However, while a television is good for displaying audiovisual information, the television doesn't do as good a job with the display of text, particularly at typical viewing distances. A computer display monitor, on the other hand, does a good job of displaying text. Thus, a computer display monitor can be used to display the secondary information. (Herein, a "computer display monitor" can display not only video, but also audio.) In particular, a portable computer (e.g., a notebook or subnotebook computer) can advantageously be used to implement such display. Moreover, the portable computer can also be used to implement the control device 101, thus allowing the display of the secondary information to be integrated with the user interface used to specify instructions for controlling operation of the system 100. Where a portable computer is used to implement the control device 101, communication between the control device 101 and the rest of the system 100 is advantageously accomplished using a wireless local area network (LAN), infrared link, or other wireless communications system, so that the user will have more freedom of movement when using the control device 101.

The system controller 103 can be implemented by any conventional processing device or devices that can accomplish the functions of a system controller as described herein. For example, the system controller 103 can be implemented by a conventional microprocessor chip, as well as peripheral and other computer chips that can be configured to perform the functions of the system controller 103. The data storage device 104 can be implemented by any conventional storage devices. The data storage device 104 can be implemented, for example, by a conventional computer hard disk (to enable storage of digital data, including analog data—e.g., television or radio signals—that has been digitized), a conventional videotape (to enable storage of, for example, analog data corresponding to acquired television signals) or a conventional audiotape (to enable storage of, for example, analog data corresponding to acquired radio signals). In particular, the system controller 103 and data storage device 104 can be implemented, for example, in a conventional digital computer. The devices with which the system controller 103 and data storage device 104 are implemented should have the capability to compress and decompress the audio, video and text data quickly enough to enable real-time display of that data. The system controller 103 can communicate with the control device 101 and the primary display device 102 in any appropriate manner, including wire and wireless communications.

In a particular embodiment of the invention, the control device 101 can be embodied by a portable computer (e.g., a Thinkpad™ computer, made by IBM Corp. of Armonk, N.Y.). The portable computer and associated display screen facilitate the presentation of a graphical user interface, as will be apparent from the description below. Preferably, the portable computer has a color display screen. A color display screen further facilitates implementation of a graphical user interface by enabling color differentiation to be used to enhance the features provided in the graphical user interface. The Thinkpad™ can be configured (as known by those skilled in such art) to act as an X/windows terminal (client) that communicates with an X/windows host (server), using standard X/windows protocols (as also known by those skilled in such art), to enable generation and display of the graphical user interface. In this particular embodiment of the

14

invention, the primary display device 102, as well as the system controller (X/windows host) 103, can be embodied, for example, by an Indigo2 workstation computer made by Silicon Graphics Incorporated (SGI) of Mountain View, Calif. The portable computer can communicate with the SGI Indigo2 computer via a wireless Ethernet link.

Alternatively, both of the primary display device 102 and control device 101 could be implemented in a digital computer with the system controller 103 and data storage device 104 (although such an implementation may not have some of the advantages of the embodiments of the invention described above). For example, the above-mentioned SGI Indigo2 computer or an IBM-compatible desktop computer could be used to implement a system of the invention in this manner. In particular, implementation of a system according to the invention in this manner could advantageously be accomplished on a portable computer such as a notebook computer.

III. User Interface

A. Graphical User Interface

1. Overview

FIG. 2A is a diagrammatic representation of a graphical user interface (GUI) 200 according to the invention that can be used to enable control of the operation of a system according to the invention, display information regarding operation of the system of the invention and display information acquired by the system of the invention. Generally, a GUI according to the invention can be displayed using any suitable display device. Further, when a GUI according to the invention is displayed on a display monitor of a digital computer, the GUI can be implemented by appropriately tailoring conventional computer display software, as known to those skilled in the art in view of the discussion below. For example, the GUI 200 can be displayed on the screen of a portable computer.

The GUI 200 includes four regions: primary information playback control region 201, primary information map region 202, related primary information region 203, and related secondary information region 204. It is to be understood that the regions 201, 202, 203 and 204 could be arranged in a different manner, have different shapes and/or occupy a greater or lesser portion of the GUI 200 than shown in FIG. 2A. Additionally, it is to be understood that a GUI according to the invention need not include all or any of the regions 201, 202, 203 or 204; it is only necessary that the GUI include features that allow the system according to the invention to be controlled. Thus, for example, a GUI according to the invention could function adequately without a related primary information region 203. The GUI also need not, for example, include a primary information map region 202 or a primary information playback control region 201 having exactly the characteristics described below; other interfaces enabling similar functionality could also be used. The GUI could also be implemented so that user interaction with standard GUI mechanisms such as menus and dialog boxes is necessary to cause display of system controls, system operation information, and/or acquired information. For example, a GUI according to the invention could be implemented such that a display of the related secondary information region 204 is produced only upon appropriate interaction with one or more menus and/or dialog boxes.

FIG. 2B is a view of an illustrative GUI 210 in accordance with the diagrammatic representation of FIG. 2A. The GUI 210 is particularly tailored for use with an embodiment of the invention in which the primary information includes

US 6,263,507 B1

15

videotape of one or more news programs and the secondary information includes the text of news stories from text news sources. Below, the regions 201, 202, 203 and 204 of the generic GUI 200 are described generally, while the corresponding regions 211, 212, 213 and 214 of the particular GUI 210 are described in detail.

2. Control of Primary Information Display

The primary information playback control region 201 of the GUI 200 is used to control the manner in which the primary information is displayed on the primary display device 102. The region 201 can be used, for example, to provide a mechanism to enable the user to begin, stop or pause display of the primary information, as well as rewind or fast forward the display. The region 201 can also be used, for example, to control the particular primary information that is displayed, as well as the apparent display rate at which the primary information is displayed.

As seen in FIG. 2B, the primary information playback control region 211 of the GUI 210 includes topic "buttons" 215, control "buttons" 216 and a speed control 217. It is to be understood that the functionality of the topic buttons 215, control buttons 216 and speed control 217, described below, could be accomplished in a manner other than that shown in FIG. 2B and described below.

The topic buttons 215 enable the user to select a subject matter category so that, for example, all news stories in the recorded news programs that pertain to the selected subject matter category are displayed one after the other by the primary display device 102. Alternatively, selection of a topic button 215 could cause a list of news stories pertaining to that subject matter category to appear, from which list the user could select one or more news stories for viewing. (The categorization of the primary information by subject matter category is discussed in more detail below.) The GUI 210 includes six topic buttons 215 to enable selection of news stories related to international news ("World"), national news ("National"), regional news ("Local"), business news ("Business"), sports news ("Sports"), and human interest news ("Living"); however, a GUI according to the invention can include any number of topic buttons and each button can correspond to any desired subject matter category designation.

The control buttons 216 enable the user to control which news story is displayed, as well as the manner in which a news story is displayed. Moving from left to right in FIG. 2B, the control buttons 216 respectively cause the display to activate a dialog box that enables the user to perform a keyword search of the text of news stories acquired by the system of the invention, return to the beginning of the currently displayed story to begin displaying the story again, stop the display, start the display, and skip ahead to the next story in a predetermined sequence of stories. A GUI according to the invention can include other control buttons that enable performance of other functions instead of, or in addition to, the functions enabled by the control buttons 216, such as fast forwarding the display, rewinding the display, pausing the display (a particular method according to the invention is described below), and displaying a summarized version of the primary information (a particular method according to the invention is described in more detail below).

The speed control 217 can be used to increase or decrease the apparent display rate with which the primary information is displayed. The speed control display 217 shows a number that represents the amount by which a normal display rate is multiplied to produce the current apparent display rate, and includes a graphical slider bar that can be used to adjust the

16

apparent display rate. The manner in which the apparent display rate can be changed is described in more detail below.

3. Map of Primary Information Display

The primary information map region 202 of the GUI 200 provides the user with a description of the content of the primary information that is available for display, as well as information that facilitates navigation through the primary information, and can also be used to allow the user to select particular primary information for display. The description of the primary information can include, for example, an illustration or other description of the subdivision of the primary information into smaller portions (e.g., segments) of information. Such illustration or description can convey the number of portions, the length (i.e., time duration) of each portion and the subject matter of each portion. The region 202 can also be used to show the user the location within the primary information of the portion of the primary information that is currently being viewed, as well as which (if any) portions of the primary information have previously been viewed. Additionally, the region 202 can be used to enable the user to move freely among portions of the primary information by, for example, using a conventional mouse to point and click on a portion of the primary information that is illustrated in the region 202.

As seen in FIG. 2B, the primary information map region 212 of the GUI 210 includes several subdivided rows, each row representing a particular news program (e.g., CNN Headline News, NBC Nightly News, etc.). Each row is a map that illustrates to some level of detail the content of the corresponding news program. Each of the subdivisions of a row represent breaks during the news program, such as breaks between news stories. The region between each subdivision represents a news story (a region could also represent, for example, an advertisement). The duration of each news story is depicted graphically by the length of the region corresponding to that news story. Each region in a row can be displayed in a particular color, each color representing a particular predetermined subject matter category (i.e., topic), so that the color of each region denotes the subject matter category of the news story corresponding to that region.

The map region 212 can be further enhanced in any of a variety of ways. For example, the news program (row) that is currently being viewed can be marked, such as by, for example, shading the row of the currently viewed news program a particular color or causing a particular type of symbol to appear adjacent to the row of the currently viewed news program. Additionally, news stories that have already been viewed can be marked in an appropriate manner, such as by, for example, causing the regions of the viewed news stories to be cross-hatched or to be shaded a particular color. The current viewing location can also be shown: in FIG. 2B, this is shown by a vertical line.

4. Related Primary Information

The related primary information region 203 of the GUI 200 displays "thumbnails" which identify segments of the primary information that are related to the primary information that is currently being displayed. Though the region 203 includes four thumbnails 203a, 203b, 203c, 203d, generally, the region 203 can be used to display any number of thumbnails. Further, the thumbnails can take any form, such as a display of a portion of the segment or a display of a representation of the segment. For example, the thumbnails 203a, 203b, 203c, 203d can be single video images that represent the video data of the segment being identified ("keyframes"). (As seen in FIG. 2B, the related primary

US 6,263,507 B1

17

information region 213 of the GUI 210 includes three single video images that each represent a news story from a news program.) Alternatively, the thumbnails 203a, 203b, 203c, 203d could be a text summary or other text identifier of the segment being identified. Or, the thumbnails 203a, 203b, 203c, 203d could be pictorial representations that identify the corresponding segment. Other possibilities exist, as known to those skilled in the art.

To enable display of thumbnails, primary information segments that are related to the primary information segment that is being displayed must be determined. A threshold of relatedness (the expression of the threshold depending upon the method used to determine relatedness) is preferably specified so that only segments that are sufficiently related to the displayed segment are displayed in the related primary information region 203, even if that means that less than the allotted number of segments (including no segments) are displayed. If appropriate, redundant segments can be eliminated from the primary information segments to be displayed in the related primary information region 203, using techniques similar to those described below for eliminating redundant segments from a set of segments identified as similar to a designated segment (e.g., eliminating redundant secondary information segments that are similar to a displayed primary information segment).

Identification of the relatedness of primary information segments can be accomplished by determining the degree of similarity between the primary information segment being displayed and each other primary information segment. The degree of similarity can be determined using any appropriate method, such as, for example, relevance feedback. The use of relevance feedback to determine the similarity between two segments is discussed in more detail below with respect to the determination of the relatedness of primary and secondary information segments (see, in particular, section IV.B.2. below). The use of relevance feedback necessitates that sets of text data that represent the primary information segments be created (by, for example, using a conventional speech recognition method to create a transcript of the spoken portion of the audio data set) if such sets of text data do not already exist (e.g., a closed-caption transcript).

When the thumbnails 203a, 203b, 203c, 203d are keyframes, each keyframe should be representative of the video content of the segment being identified. Each keyframe can be, for example, a video frame selected from the video data representing the segment. The keyframe can be selected from the video data in any appropriate manner.

For example, the keyframe can be a video frame that occurs at a specified location within the video data of the segment in a particular embodiment of the invention in which the primary information comprises television news stories, a video frame that occurs one tenth of the way through the video data representing the news story is selected. One tenth was chosen because it was determined empirically that video frames of particular relevance to the content of a television news story tend to occur at about that point in the television news story.

Alternatively, the keyframe can be selected based upon an analysis of the content of the video data. One method of accomplishing this is described in detail in the commonly owned, co-pending U.S. patent application entitled "A Method of Compressing a Plurality of Video Images for Efficiently Storing, Displaying and Searching the Plurality of Video Images," by Subutai Ahmad, U.S. Ser. No. 08/528,891, filed on Sep. 15, 1995, the disclosure of which is incorporated by reference herein. In that method, the content of each video frame is represented by a vector. The vector

18

can comprise, for example, the discrete cosine transform (DCT) coefficients for the video frame, as known to those skilled in the art of video image analysis. (The DCT coefficients indicate, for example, how much objects in a video frame have moved since the previous video frame.) From the vectors for all of the video frames of the video data of the segment an average vector is determined. The keyframe is selected as the video frame that is represented by a vector that is closest to the average vector for the video data. This method of selecting a keyframe can be advantageous as compared to the arbitrary selection of a video frame that occurs at a specified location within the video data, since it is likely to result in the selection of a video frame that is more representative of the video content of the segment.

Rather than selecting a single video frame from the video data to be the keyframe, multiple keyframes can be identified from the video data and the keyframes "tiled," i.e., presented together adjacent to each other. Or, the video data can be analyzed and a composite video frame synthesized from the video data. Any technique for synthesizing a video frame or frames can be used.

The keyframe may also be a video frame or frames that are not selected from the video data. For example, a representative video image (e.g., one or more video frames) can be selected from a library of video images. For instance, a news story about baseball could be represented by a keyframe showing a batter swinging at a pitch. Such selection can be done manually, i.e., at some point, a person reviews or is made aware of the content of the segment and, based upon that knowledge, associates a video image from the library with the segment. Alternatively, such selection can be accomplished automatically (meaning, here, without human intervention, except to establish the criteria for the selection process) by analyzing the audiovisual data of the segment (e.g., with an appropriately programmed digital computer) to ascertain the content of the segment and, based upon that analysis, associating a video image from the library with the segment. The content of the segment could be determined, for example, using a categorization method as described in more detail below. The segment to be categorized could either be compared to previously categorized segments that can be displayed by the system of the invention, or to a library of "control segments", each of which contain words germane to a particular subject.

The GUI 200 can be implemented, using conventional interface methods, so that a user of a system of the invention can select (e.g., by pointing and clicking with a mouse) one of the thumbnails 203a, 203b, 203c, 203d to cause the corresponding primary information segment to be displayed. (The map in the primary information map region 202 is adjusted accordingly.)

5. Related Secondary Information

The related secondary information region 204 of the GUI 200 provides the user information from a secondary information source or sources, the secondary information being related to the primary information currently being displayed. Though the region 204 includes two secondary information displays 204a, 204b, generally, the region 204 can include any number of secondary information displays. Further, as with the thumbnails 203a, 203b, 203c, 203d of the related primary information region 203, the secondary information displays 204a, 204b can take any form. For example, the secondary information displays 204a, 204b could be single video images, moving video images or sets of text. (As shown in FIG. 2B, the related secondary information region 214 of the GUI 210 includes three sets of text that each are a story from a text news source.) Other possibilities exist for

US 6,263,907 B1

19

the secondary information displays 204a, 204b, as known to those skilled in the art. As the segment of primary information being displayed changes, the secondary information displays 204a, 204b typically change as well. As indicated above, segments of secondary information that are related to the primary information that is being displayed can be identified in a manner discussed in more detail below. The system according to the invention can also be implemented so that the user can cause various parts of the secondary information displays 204a, 204b to be displayed, e.g., the user can be enabled to scroll up and down through a set of text or move back and forth through a video clip, using conventional GUI tools such as mouse pointing and clicking.

B. Other User Interface Techniques

User interface techniques other than GUI can be used with the invention. For example, rather than using GUI "buttons" (as illustrated in the primary information playback control region 211 of the GUI 210 of FIG. 2B), the manner in which the primary information is displayed could be controlled using a rotating knob device. Rotation of the knob in one direction could cause the display of the primary information to move forward (play); rotation of the knob in the other direction could cause the display of the primary information to move backward (rewind). Further, the knob could be constructed so that as the knob is rotated the user feels detents at certain points in the rotation. Each detent could correspond to a particular apparent display rate of the display. For example, when the knob is positioned in a home position, the display is stopped. When the knob is rotated clockwise, the display moves forward, the first detent in the clockwise direction causing the display to occur at a normal display rate, the second detent specifying a target apparent display rate of, for example, 1.5 times the normal display rate, the third detent specifying a target apparent display rate of, for example, 2.0 times the normal display rate, and so on. Similarly, when the knob is rotated counterclockwise, the display moves backward (i.e., in a chronological direction opposite that in which the display normally progresses). The first detent corresponds to normal display rate, the second detent specifies a target display rate of, for example, 1.5 times the normal display rate, and so on. The maximum rotation of the knob in either direction could be limited, the maximum rotation corresponding to a maximum target apparent display rate. The knob could be positioned at any position in between, thus allowing the target apparent display rate to be varied continuously between the maximum forward and backward display rates. The knob could also include a centrally located pushbutton to, for example, enable skipping from the display of one segment of the primary information to a next segment of the primary information. The knob could be constructed so that the position of the knob (or activation of the pushbutton) is transmitted to the remainder of the system using wireless communications, thus providing the user with relatively large freedom of movement during use of the system.

IV. Processing of Obtained Information

A. Information Acquisition

1. In General

Returning to FIG. 1, the system controller 103 causes data to be acquired from the primary information source 107 and the secondary information source 108, as described above. The data is acquired using methods and apparatus that are appropriate to the type of data being acquired. For example, the system controller 103 can acquire data representing

20

television broadcasts using conventional equipment for receiving (e.g., a television set and antenna) and recording (e.g., a conventional videocassette recorder) television signals. Or, the system controller 103 can acquire data representing radio broadcasts using conventional equipment for receiving (e.g., a radio and antenna) and recording (e.g., a conventional audiotape recorder) radio signals. Or, the system controller 103 can acquire computer-readable data files (that can include text data, audio data, video data or some combination of two or more of those types of data), using conventional communications hardware and techniques, over a computer network (e.g., a public network such as the Internet or a proprietary network such as America Online™, CompuServe™ or Prodigy™) from an information providing site that is part of that network. In one particular embodiment of the invention, the system controller 103 acquires primary information including the television signals representing the content of designated television news broadcasts, and secondary information including computer-readable data files that represent the content of designated news stories from text news sources.

The data can be acquired according to a pre-established schedule (that can be stored, for example, by the data storage device 104). Data can be acquired at any desired frequency and the scheduled acquisition times specified in any desired manner (e.g., hourly, daily at a specified time, weekly on a specified day at a specified time, or after the occurrence of a specified event). The schedule can be used, for example, to program a videocassette recorder to record particular television programs at particular times. Likewise, the schedule can be used, for example, to appropriately program a computer to retrieve desired data files from particular network sites (e.g., by specifying an appropriate network address, such as a URL) of a computer network at specified times. In the latter case, if the device with which the system controller 103 is implemented is not operating (e.g., the computer is not turned on) at a time when a scheduled acquisition of data is to take place, the system controller 103 can be implemented so that all such data is immediately retrieved upon beginning operation of the device (e.g., turning the computer on). Further, connection over the network to the site or sites from which data is to be obtained can be accomplished by, for example, inserting a communications daemon into a startup file that is executed at the beginning of operation of the operating system of a computer used to implement the system controller 103. For example, if the computer uses a Windows operating system, the daemon can initiate a Win-Sock TCP/IP connection to enable connection to be made to the network site.

The acquired data must be stored. As indicated above, analog data (such as television or radio signals) can be stored on an appropriate medium, such as videotape or audiotape. Additionally, some or all of the data acquired by a system according to the invention is, if not already in that form, converted to digital data. The digital data can be stored on a conventional hard disk having adequate capacity, as described above. To minimize the amount of data storage capacity required, the digital data can be compressed using conventional techniques and equipment. Illustratively, a half hour television news program requires approximately 250 MB of hard disk storage capacity when the video is recorded using Adobe Premiere with Radius Studio compression at 15 fps and "high" quality capture at 240x180 resolution, and the audio is recorded at approximately 22 kHz.

Appropriate rules can be established to handle situations in which the data storage device 104 (whether single or multiple devices) has insufficient data storage capacity to

US 6,263,507 B1

21

store new data. For example, the oldest data can be deleted, as necessary, to make room for new data. For example, in the particular embodiment of the invention in which the primary information is the content of designated television news programs and the secondary information is the content of designated text news stories, as new television news programs are recorded, the oldest stored programs can be deleted as necessary to make space to store the new programs, and text stories that are older than a specified length of time (e.g., several days) are automatically deleted.

The GUI 200 (FIG. 2A) can also include a mechanism for enabling the user to specify the particular information desired, i.e., specify particular information providers (e.g., news networks, such as CNN, NBC, ABC or CBS, or information services, such as Clarinet™) and data acquisition schedules for both the primary information source 107 and the secondary information source 108. This could be implemented, for example, using a set of nested menus, as known by those skilled in the art.

2. Recording/Playback Mediation

A system according to the invention may be instructed to acquire new information at the same time that the system is instructed to display other information. However, limitations of the devices or configuration of the system of the invention can impede or prevent such simultaneous acquisition and display. For example, the operating speed of a hard disk used to store the data describing the acquired information can limit the capacity of the system for such simultaneous operation: for typical amounts of audiovisual data, current conventional hard disks may not operate at a speed that is adequate to enable the simultaneous storing of data to, and accessing of stored data from, the hard disk.

Thus, in one embodiment of the invention, when data acquisition is scheduled to begin at a time when the system of the invention is being used for information display, a conventional graphical user interface mechanism (e.g., a dialog box) is used to alert the user of the system to the conflict and offer a choice between continuing with the display (thus delaying or eliminating the data acquisition) or ending the display and allowing the data acquisition to occur.

In another embodiment of the invention, the user can be alerted of an impending data acquisition at some predetermined time before the data acquisition is scheduled to begin. Similar to the choice described above, the user can be presented with a choice to continue with the display at that time or allow the data acquisition to occur. The system of the invention can default to one or the other modes of operation (i.e., data acquisition or display) if the user does not make a selection.

Or, the hard disk operating speed limitation described above can be alleviated or overcome by using multiple hard disks so that if data acquisition begins at a time when data is being accessed for use in generating a display, the newly acquired data is stored to a hard disk that does not contain any previously stored data (or that, based upon evaluation of one or more predetermined rules, does not contain data that is expected to be accessed during the time that the new data is being acquired), thus ensuring that data access and data storage will not occur simultaneously for a single hard disk. Alternatively, the hard disk operating speed limitation can be addressed by using only some portion of the available data to generate the information display, thus freeing more time for use in storing data to the hard disk. However, this latter approach may decrease the fidelity of the display unacceptably.

In a similar approach to the two hard disk approach described above, the data being acquired can be stored on a

22

data storage device of one type, while the data to be used for generating a display is accessed from a data storage device of another type. For example, incoming television signals could be stored on a videocassette tape by a VCR, while digital data from previous television transmissions is retrieved from a hard disk for use in generating a television display of the previously acquired data. The data recorded by the VCR could be digitized at a later time and stored on the hard disk for subsequent use (which use may also occur at a time at which incoming television signals are being acquired by the VCR).

B. Information Structuring

Typically, the data representing the primary and secondary information are not provided from the primary and secondary information sources in a form that enables the various aspects of the invention described herein to be realized. Thus, it is necessary or desirable to "structure" the data (i.e., to organize and categorize the data, and relate particular data to other data) in useful ways. Below are described several aspects of such data structuring that can be implemented as part of the invention.

1. Partitioning

The primary and secondary information can be, and typically are, divided ("partitioned") into smaller related sets of information of particular utility for the invention is the identification within the primary and secondary information of contiguous related sets of information that typically concern a single theme or subject and that can be delineated in some manner from adjacent information. Herein, each such contiguous related set of information can be referred to as a "segment" of the primary or secondary information. (Note that, in the description below—see section IV.C.1.—of skimming an audiovisual display, "segment" is used in a different way; there, "segment" represents a contiguous portion of a set of audio data that occurs during a specified duration of time.) Segments within the primary information are "primary information segments" while segments within the secondary information are "secondary information segments." For example, if the primary information includes the content of several news programs, the primary information can be divided into particular news programs and each news program can further be broken down into particular news stories within the news program, each news story being denoted as a segment. Similarly, if the secondary information includes content from several text sources, the secondary information can be divided into particular text sources and each text source can be further divided into separate text stories, each text story being denoted as a segment. Note that a "segment" may sometimes, strictly speaking, not be contiguous in time (though it is contiguous in content). For example, a news story that is interrupted by a commercial break, then continues after the commercial break, may be defined as a single segment, particularly if the body of information is modified so that commercial breaks—and other extraneous portions of the body of information—are eliminated (an approach that, generally, is preferred, though such portions could also be treated as segments).

Partitioning the primary and secondary information into segments is useful for a variety of reasons. For example, each segment of the primary information can be identified within the data storage device which stores the data representing the primary information, in a manner known by those skilled in the art (e.g., by maintaining a table of segment identifiers and associated locations of the beginning of the identified segment), thus enabling the primary information segments to be accessed randomly so that the user can change the displayed segment freely among the primary information

US 6,263,507 B1

23

segments. Such identification of primary information segments also enables the creation of the map region 202 of the GUI 200 (FIG. 2). Further, each segment of the primary information can be correlated, as described in more detail below, with segments of the secondary information, thereby enabling one or more secondary information segments that are sufficiently related to a primary information segment to be displayed at the same time that the primary information segment is displayed. As also described in more detail below, the correlation of primary information segments with secondary information segments can also be used to categorize the primary information segments according to subject matter, thus enabling the user to sort or to cause display of segments of the primary information that pertain to a particular subject matter category (see the discussion of the topic buttons 215 in the playback control region 211 of the GUI 210 shown in FIG. 2A).

Generally, partitioning of a set of data requires some analysis of the data to identify "breaks" within the data, i.e., differences between adjacent data that are of sufficient magnitude to indicate a significant change in the content of the information represented by the data. A break may signify a demarcation of one segment from another, but need not necessarily do so: a break may also signify, for example, a change in the video image within a segment or a change of speakers within a segment. Methods for enabling identification of breaks that constitute segment demarcation are discussed in more detail below.

Partitioning of text data is often straightforward. For example, bodies of information that are collections of segments (e.g., stories) from text sources that are represented as computer-readable data typically include markers that identify the breaks between segments. Similarly, text transcripts of bodies of information represented as a set of audiovisual information also frequently include markers that identify breaks between segments of the information. For example, closed caption text data that can accompany the audio and video data of a set of audiovisual data often includes characters that indicate breaks in the text data (most news broadcasts, for example, include closed caption text data containing markers that designate story and paragraph boundaries, the beginning and end of advertisements, and changes in speaker) and, in particular, characters that explicitly designate breaks between segments (e.g., markers that identify story boundaries). Partitioning of such text data, then, requires only the identification of the location (e.g., if the text transcript of a set of audiovisual data is time-stamped, the time of occurrence) of the markers within the text data.

Where such markers are not present, the text data can be partitioned based upon analysis of the content of the text data. In a set of audiovisual data, breaks between segments can be determined, for example, based upon identification of the occurrence of a particular word, sequence of words, or pattern of words (particularly words that typically indicate a transition), and identification of changes in speaker. As one illustration, in a news program, phrases of the form, "Jane Doe, WXYZ news, reporting live from Anytown, USA," can indicate a break between segments.

Partitioning of audio and video data typically requires some non-trivial analysis of the data. The partitioning of audio and video data in accordance with the invention can be accomplished in any suitable manner. Some examples of methods that can be used to accomplish partitioning of audio or video data are described below. (These methods are applicable to digital data; thus, if the primary information is initially analog, it must be digitized before partitioning.)

24

Typically, the audio and video data are synchronized as a result of having been recorded together. Thus, partitioning of either the audio or the video data will result in a corresponding partitioning of the other of the audio and video data. However, if the audio and video data are not synchronized, then such synchronization must be accomplished, in addition to partitioning one of the audio or video data, so that the other of the audio and video data can be partitioned in like manner.

Partitioning of audio data can be accomplished in any of a number of ways. For example, the audio data can be partitioned using a known voice recognition method. A voice recognition method that could be used with the invention is described in "A Gaussian Mixture Modeling Approach to Text-Independent Speaker Identification," by Douglas Reynolds, PhD thesis, Dept. of Electrical Engineering, Georgia Institute of Technology, 1992, the disclosure of which is incorporated by reference herein. Voice recognition methods can be tailored to, for example, identify a break in the audio data when a particular voice speaks, when a particular sequence of voices speak, or when a more complicated occurrence of voices is identified (e.g., the occurrence of two voices within a specified time of each other, or the occurrence of a voice followed by a silence of specified duration). Illustratively, when the invention is implemented as a news browser, a break between news stories could be identified when a particular newscaster's voice is followed or preceded by a silence of specified duration.

Or, the audio data can be partitioned using a known word recognition method. For example, a conventional speech recognition method (a large variety of which are known to those skilled in that art) can be used to enable identification of words. The identified words can then be analyzed in the same manner as that described above for analysis of text data, e.g., transition words or speaker changes can be used to indicate breaks. Illustratively, when the invention is implemented as a news browser, a break between news stories could be identified when one of a set of particular word patterns occurs (e.g., "we go now to", "update from", "more on that").

Audio data can also be partitioned using music recognition, i.e., a break is identified when specified music occurs. A method for partitioning audio data in this way is described in detail in the commonly owned, co-pending U.S. patent application entitled "System and Method for Selective Recording of Information," by Michelle Covell and Meg Withgott, U.S. Ser. No. 08/399,482, filed on Mar. 7, 1995, the disclosure of which is incorporated by reference herein. Partitioning of audio data using music recognition can be particularly useful when transitions between segments of the body of information are sometimes made using standard musical phrases. Illustratively, when the invention is implemented as a news browser, music recognition can be used to partition certain news programs (e.g., The MacNeil/Lehrer news hour) which use one or more standard musical phrases to transition between news stories.

Another method for partitioning audio data is pause recognition. Pause recognition is based on the assumption that a pause occurs at the time of a significant change in the content of the primary information. For many types of information, such as news programs, this is a workable assumption. A break is identified each time a pause occurs. A pause can be defined as any period of silence having greater than a specified magnitude.

Video data can be partitioned, for example, by searching for scene breaks, a method similar to the pause recognition

US 6,263,507 B1

25

method for partitioning audio data discussed immediately above. One method of accomplishing this is described in detail in the above-mentioned U.S. patent application entitled "A Method of Compressing a Plurality of Video Images for Efficiently Storing, Displaying and Searching the Plurality of Video Images," by Subutai Ahmad. In that method, the content of each video frame is represented by a vector, as described above. The vector for each video frame is compared to the vector of the immediately previous video frame and the immediately subsequent video frame, i.e., vectors of adjacent video frames are compared. In one approach, a break is identified each time the difference between the vectors of adjacent video frames is greater than a predetermined threshold. In another approach, a predetermined number of partitions is specified and the video frames are partitioned to produce that number of partitions (the partitioning can be accomplished by considering each video frame to be initially partitioned from all other video frames and recursively eliminating the partition between partitioned video frames having the least difference, or considering none of the video frames to be partitioned and recursively establishing partitions between unpartitioned video frames having the greatest difference).

Other approaches to scene break identification could be used, as known by those skilled in the art of processing video images. Some other approaches to scene break identification are discussed in "Automatic Parsing of News Video," by Hongliang Zhang, Gong Yihong, Stephen W. Smoliar, and Tan Ching Yong, IEEE Conference on Multimedia Computing and Systems, Boston, May 1994, the disclosure of which is incorporated by reference herein. For example, scene breaks could be identified based upon the magnitude of the overall changes in color of the pixels of adjacent video frames (a color change having a magnitude above a specified threshold is identified as a scene break). Or, scene breaks could be identified based upon the magnitude of the compression ratio for a particular set of adjacent video frames (a relatively small amount of compression indicates a relatively large change between video frames and, likely, a change in scenes, i.e., a scene break).

The above-described methods for partitioning audio or video data directly may not, by themselves, enable identification of segment breaks to be accomplished easily or at all. For example, without augmentation, pause recognition or scene break identification typically are not implemented in a manner that enables distinguishing between segment breaks and other breaks. Voice recognition may not, alone, be a reliable indicator of segment breaks, since switches in speaker often occur for reasons unrelated to a segment break. Word recognition, too, may be erratic in determining segment breaks; it also requires obtaining a text transcript of the audio. Music recognition works well only with a limited number of information sources, i.e., information sources that use well-defined musical transitions.

It may be possible to include markers (similar to those discussed above with respect to closed caption text data) in either audio or video data that directly identify segment or other breaks within the audio or video data. The invention contemplates use of such markers to segment audio and/or video data.

If a set of audiovisual data also includes text data (e.g., a closed caption transcript of the spoken audio), it is possible to partition the audiovisual data by partitioning the text data, then using the partitioned text data to partition the audio data and video data in a corresponding manner. Even if the audiovisual data does not initially include text data, the text data can be produced using a speech recognition method.

26

The text data can be partitioned using any appropriate method, as described above.

Typically, the text data, audio data and video data are each time-stamped. Theoretically, then, once segment breaks are determined in the text data, the time-stamps of the beginning and end of each segment within the text data could be used directly to identify segment breaks within the audio data and/or video data. However, in practice, the text data is typically not exactly synchronized with the audio data and video data (e.g., the text data of a particular segment may begin or end several seconds after the corresponding audio or video data), making such a straightforward approach infeasible. Nevertheless, the time-stamps of the segment breaks in the text data can be used to enable synchronization of those segment breaks with the corresponding segment breaks in the audio and video data. Such synchronization can be accomplished using any appropriate technique. Some possible approaches are described below.

One way to partition the audio and video data based upon the partition of the text data is to use a synchronization of the complete set of audio data with the complete set of text data, and a synchronization of the complete set of audio data with the complete set of video data to identify the partitions in the audio and video data. The latter synchronization typically exists as a consequence of the manner in which the audio and video data is obtained. However, synchronization between the text data and the audio data frequently does not already exist, and, if it does not, obtaining such synchronization can be computationally expensive. Further, it is not necessary to synchronize all of the text data with the audio and video data, but, rather, only the locations of the segment breaks.

A simpler approach is to determine the segment breaks in the audio and video data from the segment breaks in the text data based upon a rule or rules that exploit one or more characteristics of the body of information. Such a rule might be based on an observation that segment breaks in the audio and/or video data of a set of audiovisual data bear a relatively fixed relationship to the corresponding segment breaks in the corresponding text data. For example, it was observed that the video data of a news story from an audiovisual news program frequently begins about 5 to 10 seconds before the closed caption text data of the news story. Thus, in one embodiment of news browser implementation of the invention, the beginning of the video data of a news story is assumed to be 4 seconds prior to the closed-caption text data. This enables most of the relevant video data to be captured, while reducing the possibility of capturing extraneous video. This approach was found to be accurate within 2 seconds for CNN Headline News and the news programs of the NBC, ABC and CBS television broadcasting networks.

In some cases, the approach may still not produce as good a result as desired, i.e., the segmentation of the audio and video data is not as crisp as desired, either deleting part of the beginning or end of the audio or video segment, or including extraneous audio or video as part of the segment. Thus, according to another particular embodiment of the invention, partitioning of audiovisual data that includes text data in which segments breaks are explicitly designated by markers within the text data can be accomplished in two steps: a first, coarse partitioning followed by a second, fine partitioning. FIG. 3 is a flow chart of a method 300, in accordance with this aspect of the invention, for identifying the boundaries of segments in a body of information. In the coarse partitioning step 301 of the method 300, the time-stamps associated with the segment breaks in the text data

US 6,263,507 B1

27

can be used to approximate the location of the corresponding segment breaks in the audio and video data, as described above. In step 302, a window of data (e.g., audio or video data in the context of the current discussion) that includes the approximate segment boundary is specified. This can be accomplished, for example, by specifying a time range that includes the time associated with the segment break in the text data (e.g., the time of occurrence of the segment break in the text data plus or minus several seconds) and identifying audio and/or video data that falls within that time range from the time-stamps associated with the audio and/or video data. The fine partitioning step 303 can then be used to identify breaks within the audio and/or video data. The fine partitioning can be accomplished using any appropriate method, such as one of the above-discussed methods (i.e., scene break identification, pause recognition, voice recognition, word recognition, or music recognition) to identify breaks in audio and video data. The fine partitioning can be performed on the entire set of audio data or video data, or only on the audio or video data that occurs within the time range. In the step 304, the data within the time range can then be examined to identify the location of a break or breaks within the time range. If more than one break is identified, the "best" break, measured according to the criteria of the partitioning method used, can be identified as the segment break, or the break occurring closest in time to the approximate segment break can be identified as the segment break.

Once the segment breaks in the audio or video data are identified, segment breaks in the other of the audio or video data can be determined using a synchronization of the audio and video data, as discussed above. Pointers to the segment breaks in the text data, audio data and/or video data can be maintained to indicate the beginning and end of each segment, thus enabling random access to segments within a body of information (e.g., news stories within a news program), as discussed in more detail above. The identified segments can also be used to enable other features of the invention, as described in more detail below.

2. Correlation

As mentioned above, the related secondary information region 204 of the GUI 200 is used to provide the user, from a secondary information source or sources, information that is related to the primary information currently being displayed. Thus, it is necessary to determine which of the segments of the secondary information are sufficiently related to the primary information segment displayed on the primary display device 102 to be displayed in the related secondary information region 204. This can be accomplished by determining the degree of similarity between each segment of the primary information (e.g., news story from an audiovisual news program) and each segment of the secondary information (e.g., text story from a text news source), and displaying in the related secondary information region 204 of the GUI 200 certain secondary information segments that are most similar to the primary information segment that is being displayed by the primary display device 102.

An important aspect of the invention is the capability to determine relatedness of segments of information represented by two different types of data. In particular, the invention can enable the determination of relatedness between segments of information represented by audiovisual data (such as is frequently the case for the primary information that can be displayed by the invention) and segments represented by text data (such as is generally the case for the secondary information as described particularly

28

herein). This aspect of the invention enables the display of the related secondary information region 204 to be generated. It can also enable categorization of uncategorized segments, as described further below.

FIG. 4 is a flow chart of a method 400, in accordance with this aspect of the invention, for determining whether a first set of information represented by a first set of data of a first type (e.g., audiovisual data) is relevant to a second set of information represented by a second set of data of a second type (e.g., text data). In step 401, a set of data of the second type is derived from the first set of data of the first type. In a typical application of the method 400, step 401 causes a set of text data to be produced from a set of audiovisual data. The set of text data can be produced in any appropriate manner. For example, "production" of the set of text data may be as simple as extracting a pre-existing text transcript (e.g., a closed caption transcript) from the set of audiovisual data. Or, the set of text data can be produced from the set of audio data using a conventional speech recognition method. In step 402, the derived set of data (of the second type) is compared to the second set of data of the second type to determine the degree of similarity between the derived set of data and the second set of data. One way of making this determination is described in more detail below. In step 403, a determination is made as to whether the first set of data is relevant to the second set of data, based on the comparison of step 402. Typically, a threshold level of similarity (the expression of which depends upon the method used to determine similarity) is specified so that only sets of information that are sufficiently related to each other are identified as related. (This means, when the method 400 is used to generate the related secondary information region 204, that less than the allotted number of secondary information segments—or even no secondary information segments—may be displayed.)

The degree of similarity can be determined using any appropriate method, such as, for example, relevance feedback. In relevance feedback, a text representation of each segment to be compared (e.g., each audiovisual news story or text story) is represented as a vector, each component of the vector corresponding to a word, the value of each component being the number of occurrences of the word in the segment. (Two words are considered identical—i.e., are amalgamated for purposes of ascribing a magnitude to each component of the vector representing the textual content of a segment—if the words have the same stem; for example, "play", "played" and "player" are all considered to be the same word for purposes of forming the segment vector.) For each pair of segments, the normalized dot product of the vectors corresponding to the segments is calculated, yielding a number between 0 and 1. The degree of similarity between two segments is represented by the magnitude of the normalized dot product, 1 representing two segments with identical words and 0 representing two segments having no matching words. The use of relevance feedback to determine the similarity between two text segments is well-known, and is described in more detail in, for example, the textbook entitled *Introduction to Modern Information Retrieval*, by Gerard Salton, McGraw-Hill, New York, 1983, the pertinent disclosure of which is incorporated by reference herein. Relevance feedback is also described in detail in "Improving Retrieval Performance by Relevance Feedback," Salton, G., *Journal of the American Society for Information Science*, vol. 41, no. 4, pp. 288-297, June 1990 as well as "The Effect of Adding Relevance Information in a Relevance Feedback Environment," Buckley, C. et. al., *Proceedings of 17th International Conference on Research and Development in*

US 6,263,507 B1

29

Information Retrieval, DIGIR 94, Springer-verlag (Germany), 1994, pp. 292-300, the disclosures of which are incorporated by reference herein.

The related secondary information region 204 of the GUI 200 can display a predetermined number of relevant secondary information segments. Generally, it is desirable to display the secondary information segments that are most similar to the primary information segment that is being displayed. While this can be accomplished straightforwardly by displaying those secondary information segments having the highest determined degree of similarity, such an approach may not be desirable in some situations. For example, the secondary information source may include segments that are identical or nearly identical (e.g., news stories are often repeated in a variety of text news sources with little or no change), so that display of the secondary information segments having the highest determined degree of similarity can result in undesirable redundancy.

This problem can be overcome by further determining the degree of similarity between each of a predetermined number of the secondary information segments having the highest determined degree of similarity (in one embodiment of the news browser implementation of the invention, the 10 most similar text stories are compared), and displaying only one of each pair of secondary information segments having a degree of similarity above a specified threshold, i.e., redundant secondary information segments are eliminated. Again, this can be more problematic than first appears. For example, a particular segment may have greater than the threshold degree of similarity when compared to each of second and third segments, but the second and third segments may have less than the threshold degree of similarity when compared to each other. From the three segments, it would be desirable to show both the second and third segments. However, if the first segment is compared to the second segment or the third segment, and the second or third segment discarded, before comparison of the first segment to the other of the second or third segment (which will also result in discarding of one of the compared segments), then only one of the three segments will be shown. Such a situation could be handled by, for example, calculating the similarity between all pairs of the predetermined number of secondary information segments, and performing comparisons that reveal the situation described above before discarding any of the secondary information segments.

3. Categorizing

An important aspect of the invention is the capability to categorize uncategorized segments of information based upon the categorization of previously categorized segments of information. In particular, if the segments of the secondary information have been categorized according to subject matter, then the degree of similarity between the subject matter content of segments of the primary information (e.g., news stories in audiovisual news programs) and segments of the secondary information (e.g., news stories from text news sources) can also be used to categorize the primary information according to subject matter. This can be useful to enable determination of which primary information segments fall within a particular subject matter category that corresponds to one of the topic buttons 215 (FIG. 2) that a user can select to cause all primary information segments that pertain to the selected subject matter category to be displayed one after the other by the primary display device 103 (FIG. 1). Though this aspect of the invention has particular utility in categorizing primary information segments based upon the categorization of pre-existing secondary information segments, it can generally enable any categorized segments to be used to categorize uncategorized segments.

30

FIG. 3 is a flow chart of a method 300, in accordance with this aspect of the invention, for categorizing according to subject matter an uncategorized segment of a body of information based on the subject matter categorization of other previously categorized segments of the body of information. For example, each story from the Clarinet™ news service is categorized according to the subject matter of the story by associating one or more predefined subject matter categories (e.g., sports, travel, computers, business, international news) with the story. This subject matter categorization can be used to categorize news stories from audiovisual news programs based on the similarity between each audiovisual news story and text stories from the Clarinet™ news service. Below, such categorization of audiovisual news stories is described as an example of how categorizing segments of primary information can be accomplished in accordance with the invention.

The subject matter category or categories associated with each Clarinet™ text story are acquired as part of the acquisition of the text stories themselves and can, for example, be stored in a relational database in a memory that is part of the system controller 103 (FIG. 1). It may be desirable to associate only one subject matter category with each text story. For example, the most salient subject matter category can be identified in any appropriate manner and used as the sole subject matter category associated with the story. This may be done, for example, to increase the likelihood that the subject matter category eventually associated with each news story accurately describes the subject matter content of that news story.

In step 301 of the method 300, a determination is made as to the degree of similarity between the subject matter content of an uncategorized segment and that of previously categorized segments. The degree of similarity can be determined using any appropriate method, such as, for example, relevance feedback. When relevance feedback is used, it is necessary to obtain a textual representation of audiovisual data, if appropriate (i.e., if one or both of the segments is represented as audiovisual data) and not already existent.

In step 302, previously categorized segments that are relevant to the uncategorized segment are identified. Relevant segments can be identified based upon the degree of similarity in the same manner as that described above with respect to correlation of segments, e.g., segments having greater than a threshold level of similarity can be designated as relevant. Step 301 can also include elimination of redundant segments (in the same manner as described above) from among those that have the required degree of similarity to the uncategorized segment.

In step 303, the uncategorized segment is categorized based upon the subject matter categories associated with the relevant previously categorized segments. One or more subject matter categories can be associated with the uncategorized segment. Generally, the subject matter category or categories can be selected from the subject matter categories associated with the relevant previously categorized segments using any desired method. For example, the subject matter category or categories of the most similar previously categorized segment could be selected as the subject matter category or categories of the uncategorized segment. Or, the most frequently occurring subject matter category or categories associated with a predefined number of the most similar previously categorized segments (or previously categorized segments having greater than a threshold degree of similarity) could be selected as the subject matter category of the uncategorized segment. In the latter case, it may be particularly desirable, as described above, to determine the

US 6,263,507 B1

31

similarity between the relevant previously categorized segments, so that only one of a set of previously categorized segments that are substantially identical to each other influences the categorization of the uncategorized segment.

C. Information Presentation

Above, the acquisition of information and the structuring of acquired information has been described. The information must, of course, also be displayed to a user. The information display has been described generally above with respect to FIGS. 2A and 2B. However, a system according to the invention can also include one or more of a variety of additional features that enhance the information display.

1. Skimming

As indicated above with respect to FIGS. 2A and 2B, the apparent display rate with which the primary information is displayed by the primary display device 102 can be varied by the user. Variation in the apparent display rate of an audiovisual display can be implemented by appropriately programming a digital computer to accomplish the functions of a method for varying the apparent display rate. Generally, any method for varying the apparent display rate can be used with the invention. As described elsewhere herein, the primary information will often be represented by coextensive sets of data of several types (audio, video and, possible text). The particular method used to vary the apparent display rate of the primary information will typically depend upon the type of the set of data (e.g., audio, video, text) that is directly modified to produce appropriately modified data for use in generating a display of the primary information at the new apparent display rate. The method also preferably synchronizes the sets of data that are not directly modified with the set of data that is.

For example, the audio data can be modified to cause the apparent display rate of the audio display to be varied (either slowed down or speeded up) from a normal display rate and the video data synchronized with the modified audio data (resulting in a variation of the apparent video display rate that corresponds to the variation in the apparent audio display rate). Several methods of accomplishing such variation in the apparent display rate of an audiovisual display are described in detail in the commonly owned, co-pending U.S. patent application entitled "Variable Rate Video Playback with Synchronized Audio," by Neal A. Bhadkamkar, Subutai Ahmad and Michelle Covell, attorney docket number 10359-991168, filed on the same day as the present application, the disclosure of which is incorporated by reference herein. At least some of the methods described therein have the advantage that the apparent display rate of the audio can be varied while maintaining proper pitch (i.e., the voices don't sound stupefied when the display is slowed down or like chipmunks when the display is speeded up) and, therefore, intelligibility. A brief description of a general method described therein is given immediately below, followed by a brief description of one particular method for modifying the audio data.

Generally, in the methods described in the above-mentioned patent application, a correspondence between an original audio data set and an original video data set is first established. For example, the number of audio samples that have the same duration as a frame of video data can be determined and that number of audio samples defined to be an audio segment. (Note that, as mentioned above, as used here in the description of skimming, "segment" refers to a contiguous portion of a set of audio data that occurs during a specified duration of time; elsewhere herein, "segment" refers to a contiguous related set of information within the primary or secondary information that typically concerns a

32

single theme or subject and that can be delineated in some manner from adjacent information.) The audio segments can be defined, for example, so that each audio segment corresponds to a single particular video frame. A target display rate (which can be faster or slower than a normal display rate at which an audiovisual display system generates an audiovisual display from the unmodified, original sets of audio and video data) is also determined. The target display rate can be a single value which remains unchanged throughout the display or a sequence of values such that the target display rate changes during the display. The original audio data set is manipulated, based upon the target display rate and an evaluation of the original audio data set, to produce a modified audio data set. As described below, the modified audio data set is produced so that, generally, when the modified audio data set is used to generate an audio display, the audio display appears to be speeded up or slowed down by an amount that is approximately equal to the target display rate. The correspondence between the modified audio data set and the original audio data set, and the correspondence between the original audio data set and the original video data set, are used to create a correspondence between the modified audio data set and the original video data set, which, in turn, is used to delete video data from, or add video data to, as appropriate, the original video data set to create a modified video data set. Once the modified audio and video data sets have been created, an audiovisual display can be generated from those modified data sets by an audiovisual display system, or the modified audio and video data sets can be stored on a conventional data storage device for use in generating a display at a later time. The audio and video data of the modified audio and video data sets are processed at the same rate as before (i.e., when the original audio and video data sets were used to generate a display at the normal display rate) by the audiovisual display system. However, since the modified audio and video data sets (in the usual case) have a different amount (either more or less) of data than the original audio and video data sets, the apparent display rate of the audiovisual display generated from the modified audio and video data sets is different than the normal display rate. Further, since the modified video data set is created based upon the content of the modified audio data set and a correspondence between the modified audio data set and the original video data set, the modified video data set is synchronized (at least approximately and, possibly, exactly) with the modified audio data set and produces a display of the same or approximately the same duration.

The audio data can be modified in any suitable manner; one way is described following. An audio data set is divided into non-overlapping segments of equal length. Generally, the beginning and end of each segment are overlapped with the end and beginning, respectively, of adjacent segments. (Note that the overlap can be negative, such that the length of the adjacent segments is extended. The audio data of corresponding overlapped portions of adjacent segments are blended and replaced by the blended audio data. The possible lengths of each overlap are constrained in accordance with a target overlap that corresponds to the specified target display rate. However, within this constraint, the length of each particular overlap is chosen so that the pitch pulses of the overlapped portions closely resemble each other. Consequently, the blending of the audio data of the overlapped portions does not greatly distort the sound corresponding to the overlapped portions of audio data. Thus, the invention enables the audio data set to be condensed or expanded a desired amount (i.e., the display of an audio data

US 6,263,507 B1

33

set can be speeded up or slowed down as desired), while minimizing the amount of distortion associated with the modification of the audio data set (i.e., the audio display sounds "normal").

Since the actual amount of overlap of segments can vary from the target overlap that corresponds to the specified target display rate, the actual apparent display rate can vary from the target display rate. Over relatively long periods of time (e.g., greater than approximately 0.5 seconds), the actual apparent display rate typically closely approximates the target display rate. Over shorter time periods (e.g., approximately 30 milliseconds), the actual apparent display rate can vary more substantially from the target display rate. However, these short term fluctuations are not perceptible to an observer. Thus, this method produces an actual apparent display rate that to an observer appears to faithfully track the target display rate over the entire range of the display.

Preferably, the computation required to produce a particular amount of variation in the apparent display rate is done at the time that the determination of a target display rate mandates such variation. This has the advantage of reducing the amount of data storage capacity required by a system of the invention. This also enables any magnitude of apparent display rate to be specified over a continuous range of allowed display rates, rather than restricting the magnitude of the apparent display rate to one of a set of discrete magnitudes within an allowed range, as would be necessary if all of the computations for each magnitude of apparent display rate were pre-computed. Additionally, this enables the apparent display rate of the display to be varied in real time.

2. Summarization

A system according to the invention can include another information presentation feature that enables the display of a primary segment or segments to be summarized. Summarization enables an observer to quickly get an overview of the content of a particular segment or segments of information. Summarization can be implemented by appropriately programming a digital computer to accomplish the functions of a summarization method. Generally, summarization can be accomplished using any appropriate method. As with skimming, discussed above, the particular method used will typically depend upon the type of the set of data (e.g., audio, video, text) that is directly modified to produce appropriately modified data for use in generating a summary display of the primary information. The method also preferably synchronizes the sets of data that are not modified directly with the set of data that is.

For example, text data that is part of, or derived from, audiovisual data that represents a primary segment can be summarized, and the corresponding audio and video data summarized based upon the text summary. One method of accomplishing such summarization is described in detail in the commonly owned, co-pending U.S. patent application entitled "Indirect Manipulation Of Data Using Temporally Related Data, With Particular Application To Manipulation Of Audio Or Audiovisual Data," by Emanuel E. Farber and Subutai Ahmad, attorney docket number K0359-991110, filed on the same day as the present application, the disclosure of which is incorporated by reference herein. A brief description of that method is given immediately below.

The text data of a set of audiovisual data represents a transcription of the spoken portion of the audio data and is temporally related to each of the audio and video data. The text data can be obtained in any appropriate manner, e.g., the text data can be pre-existing text data such as closed-caption data or subtitles, or the text data can be obtained by using

34

any of a number of known speech recognition methods to analyze the audio data to produce the text data.

The text data is summarized using an appropriate summarization method. Generally, any text summarization method can be used; a particular example of a text summarization method that can be used with the invention is described in U.S. Pat. No. 5,384,703, issued to Withgott et al. on Jan. 24, 1995.

The unsummarized text data is aligned with the unsummarized audio data. If the text data has been obtained from the audio data using a speech recognition method, then the alignment of the unsummarized text data with the unsummarized audio data typically exists as a byproduct of the speech recognition method. Otherwise, alignment is accomplished in three steps. First, the unsummarized text data is evaluated to generate a corresponding linguistic transcription network (e.g., a network describing the set of possible phonetic transcriptions). Second, a feature analysis is performed on the audio samples comprising the unsummarized audio data set to create a set of audio feature data. Third, the linguistic transcription network is compared to the set of audio feature data (using Hidden Markov Models to describe the linguistic units of the linguistic transcription network in terms of audio features) to determine the linguistic transcription (from all of the possible linguistic transcriptions allowed by the linguistic transcription network) which best fits the set of audio feature data. As a result of this comparison, the audio features of the best fit linguistic transcription are correlated with audio features in the set of audio feature data. The audio features of the best fit linguistic transcription can also be correlated with the linguistic units of the linguistic transcription network. The linguistic units of the linguistic transcription network can, in turn, be correlated with the unsummarized text data. As a consequence of these correlations, an alignment of the unsummarized text data with the unsummarized audio data can be obtained. Using the previously determined text summary and the alignment between the text data and audio data, an audio summary can be produced.

A video summary can be produced from the audio summary using an alignment between the unsummarized audio data and the unsummarized video data. Such alignment can be pre-existing (because the audio data and video data were recorded together, the alignment being inherent because of the like time stamps associated with each of the audio and video data) or can be calculated easily (the time stamp for an audio sample or video frame can be calculated by multiplying the time duration of each sample or frame by the sequence number of the sample or frame within the audio data or video data).

Another method that can be used to summarize the display of a set of audiovisual information includes identifying and eliminating "sound bites" (defined below) in the audio portion of the primary information. The sound bites can be identified based upon analysis of a set of text data that corresponds to the spoken portion of the set of audio data. The text data can be obtained in any appropriate manner. For example, the text data may be closed caption data that is provided with the audio and video data representing the primary information. Or, the text data can be obtained from the set of audio data using conventional speech recognition techniques. Once the text data is obtained, the text data can be "pre-processed" using known methods to classify the words in the text data according to their characteristics, e.g., part of speech.

Herein, a "sound bite" is a related set of contiguous audio information that conforms to one or more predetermined

US 6,263,507 B1

35

criteria that are intended to identify short spoken phrases that are not spoken by a previously identified primary speaker and that represent information of little interest and/or are redundant. For example, in a news browser according to the invention, where the primary information includes the content of audiovisual news programs (e.g., television news programs), the predetermined criteria can be established so that spoken portions of the audio information that are likely not to have been spoken by a news anchorperson or a news reporter are identified as sound bites. Such criteria might include, for example, rules that tend to identify a spoken portion of the audio as a sound bite if the spoken portion includes slang words or the use of first person pronouns (e.g., I or we), both of which tend not to be present in the speech of an anchorperson or reporter. As can be appreciated, elimination of such audio portions will typically not significantly adversely affect the presentation of the essential content of a set of audio information, but will enable the set of audio information to be presented more quickly. (It should be noted that the summarization method of Withgott et al. was also found to be incidentally effective at eliminating sound bites.)

Once the audio data has been modified by eliminating the audio data corresponding to the sound bites, the set of modified audio data must be aligned (synchronized) with the video data (if present) to enable the video data to be modified to produce a speeded-up video display. As described above with respect to the summarization method of Farber and Ahmad, the audio/video alignment can either be pre-existing or calculated easily.

As can be appreciated, a summarization method such as one of those described above could be used in combination with a method for increasing the apparent display rate as described above (see section IV.C.1. above on Skimming) to even further condense the display of a set of primary information. For example, the set or sets of data representing the primary information could be modified to increase the apparent display rate, then the modified set or sets of data could be summarized to produce a speeded-up summary of the set of primary information. Or, conversely, the set or sets of data representing the primary information could be summarized, then the summarized set or sets of data modified to increase the apparent display rate, thus producing a speeded-up summary of the set of primary information.

As can be appreciated, the methods described above for manipulating audiovisual data to produce a summarized display of the audiovisual data can also be used, with appropriate modification (e.g., instead of producing a summary of the text data, the text data could be manipulated in some other desired fashion), to manipulate the audiovisual data for some other purpose, such as rearranging, editing, selectively accessing or searching the audiovisual data.

3. Display Pause with Elastic Playback

A system according to the invention can include yet another information presentation feature that enables the display of an image to be paused, then, at the end of the pause, resumed at an accelerated rate (i.e., a rate that is faster than a normal display rate) until a time at which the content of the display corresponds to the content that would have been displayed had the image been displayed at the normal display rate without the pause, at which time display of the image at the normal display rate resumes. In other words, after a pause, the image display is speeded up so that the display "catches up" to where it would have been without the pause, then slowed back down to the normal display rate. The implementation of this feature is described in detail in the commonly owned, co-pending U.S. patent application

36

entitled "Display Pause with Elastic Playback," by Subutai Ahmad, Neal A. Bhadkamkar, Steve B. Cousins, Paul A. Freiburger and Brygg A. Ullmer, attorney docket number I0359-991150, filed on the same day as the present application, the disclosure of which is incorporated by reference herein. A brief description of the implementation is given immediately below.

The image to be displayed is represented by an ordered set of display data. This display data is acquired from a data source at a first rate. The display data is transferred to a display device at the first rate as the display data is acquired. An image is generated from the display data transferred to the display device and displayed on the display device. At some point, the user instructs the system to pause the display. The system identifies the pause instruction from the user and, in response, stops the transfer of display data to the display device and begins storing the acquired display data at the first rate. At some later time, the user instructs the system to resume the display. The system identifies the resume instruction from the user and, in response, begins transferring stored display data to the display device at a second, effective rate that is greater than the first rate. An image is generated from the stored display data transferred to the display device and displayed on the display device. While the stored display data is being transferred to the display device, the newly acquired data continues to be stored. The storage of display data finally stops when there is no more stored display data to be transferred to the display device, the amount of stored display data having gradually been reduced by transferral of the stored display data to the display device at the second, effective rate that is greater than the first rate at which the display data is stored. Once the storage of display data stops, the display data is again transferred to the display device at the first rate as the display data is acquired.

This feature of the invention enables a great deal of flexibility in observing a real-time display of audiovisual information. For example, the invention enables an observer to pause and resume the display as desired so that, if the observer wants to temporarily stop watching to go to the bathroom or to take a phone call, the observer can pause the display, then, after resuming the display upon return, watch the audiovisual information at an accelerated display rate until the display of the program catches up to where it would have been without the pause. Thus, the user can attend to other matters while the audiovisual information is being viewed, without sacrificing viewing any of the content of the audiovisual information or enduring the inconvenience of spending additional time to finish watching the audiovisual program. This feature of the invention can also be tailored to enable a user who has begun viewing the audiovisual information at a time later than desired, to observe the audiovisual information at an accelerated rate until the display catches up to the point at which the display have been if the audiovisual information had been viewed at a normal display rate beginning at the desired start time.

Various embodiments of the invention have been described. The descriptions are intended to be illustrative, not limitative. Thus, it will be apparent to one skilled in the art that certain modifications may be made to the invention as described without departing from the scope of the claims set out below.

We claim:

1. A system for acquiring and reviewing a body of information, wherein the body of information includes a plurality of segments, each segment representing a defined set of information in the body of information, the system comprising:

US 6,263,507 B1

37

means for acquiring data representing the body of information;

means for storing the acquired data;

first display means for generating a display of a first segment of the body of information from data that is part of the stored data;

means for comparing data representing a segment of the body of information to data representing a different segment of the body of information to determine whether, according to one or more predetermined criteria, the compared segments are related; and

second display means for generating a display of a portion of, or a representation of, a second segment of the body of information from data that is part of the stored data, wherein the second display means displays the portion or representation of the second segment in response to the display by the first display means of a first segment to which the second segment is related.

2. A system as in claim 1, wherein the second display means displays the portion or representation of the second segment substantially coextensive in time with the display of the related first segment by the first display means.

3. A system as in claim 1, wherein:

at least a portion of the body of information is represented by audiovisual data;

the first segment is represented by audiovisual data;

the first display means displays an audiovisual display of the first segment; and

the second segment is represented by audiovisual data.

4. A system as in claim 3, further comprising means for selecting a segment for which a portion or representation is displayed by the second display means, wherein selection of such segment causes the first display means to display an audiovisual display of the selected segment.

5. A system as in claim 1, wherein:

at least a portion of the body of information is represented by audiovisual data;

the first display means displays an audiovisual display of the first segment; and

the second display means displays a text display of a portion or representation of the second segment.

6. A system as in claim 1, wherein:

the first display means is an analog display device; and the second display means is a digital display device.

7. A system as in claim 1, wherein:

the first display means is a television; and

the second display means is a computer display monitor.

8. A system as in claim 1, further comprising means for identifying the subject matter content of a segment of the body of information, wherein the means for comparing further comprises means for determining the similarity of the subject matter content of a segment to the subject matter content of a different segment, the predetermined criteria including a predefined degree of similarity with respect to which the relatedness of the compared segments is determined.

9. A system as in claim 8, wherein the means for determining the similarity of the subject matter of segments further comprises means for performing a relevance feedback method.

10. A system as in claim 1, wherein the means for acquiring data further comprises means for acquiring television broadcast signals.

11. A system as in claim 1, wherein the means for acquiring data further comprises means for acquiring radio broadcast signals.

38

12. A system as in claim 1, wherein the means for acquiring data further comprises means for acquiring computer-readable data files over a computer network from an information providing site that is part of that network.

13. A system as in claim 1, wherein the means for acquiring data further comprises:

means for acquiring television broadcast signals; and

means for acquiring computer-readable data files over a computer network from an information providing site that is part of that network.

14. A system as in claim 13, wherein:

the first segment is represented by data produced from the television broadcast signals; and

the second segment is represented by data from the computer-readable data files.

15. A system as in claim 1, further comprising means for identifying an instruction from a user to begin displaying at least some of the body of information, wherein the first display means begins displaying a segment in response to the user instruction.

16. A system as in claim 1, wherein the first and second display means are physically separate.

17. A system as in claim 1, wherein the means for storing the acquired data, the first display means and the second display means are interconnected to a conventional computer bus that enables the devices to communicate with each other such that the devices do not require wire communication over network communication lines to communicate with each other.

18. A system as in claim 1, wherein at least some of the acquired data is digital data, the means for acquiring data further comprising means for acquiring digital data.

19. A system as in claim 1, wherein at least some of the acquired data is analog data, the means for acquiring data further comprising means for acquiring analog data.

20. A method for acquiring and reviewing a body of information, wherein the body of information includes a plurality of segments, each segment representing a defined set of information in the body of information, the method comprising the steps of:

acquiring data representing the body of information;

storing the acquired data;

generating a display of a first segment of the body of information from data that is part of the stored data;

comparing data representing a segment of the body of information to data representing a different segment of the body of information to determine whether, according to one or more predetermined criteria, the compared segments are related; and

generating a display of a portion of, or a representation of, a second segment of the body of information from data that is part of the stored data, wherein the display of the portion or representation of the second segment is generated in response to the display of a first segment to which the second segment is related.

21. A method as in claim 20, further comprising the step of causing the display of the portion or representation of the second segment to occur substantially coextensive in time with the display of the related first segment.

22. A method as in claim 20, wherein:

the step of acquiring data representing the body of information further comprises the step of acquiring audiovisual data representing at least a portion of the body of information, wherein the first and second segments are represented by audiovisual data; and

US 6,263,507 B1

39

the step of generating a display of a first segment of the body of information further comprises the step of generating an audiovisual display of the first segment.

23. A method as in claim 22, further comprising the step of identifying the selection of a second segment for which a portion or representation is being displayed, wherein selection of such second segment causes an audiovisual display of the selected second segment to be produced.

24. A method as in claim 20, wherein:

the step of acquiring data representing the body of information further comprises the step of acquiring audiovisual data representing at least a portion of the body of information;

the step of generating a display of a first segment of the body of information further comprises the step of generating an audiovisual display of the first segment; and

the step of generating a display of a portion of, or a representation of, a second segment of the body of information further comprises the step of generating a text display of the portion or representation of the second segment.

25. A method as in claim 20, wherein:

the step of generating a display of a first segment of the body of information further comprises the step of generating a display of the first segment on an analog display device; and

the step of generating a display of a portion of, or a representation of, a second segment of the body of information further comprises the step of generating a display of the portion or representation of the second segment on a digital display device.

26. A method as in claim 20, wherein:

the step of generating a display of the first segment on an analog display device further comprises the step of generating a display of the first segment on a television; and

the step of generating a display of the portion or representation of the second segment on a digital display device further comprises the step of generating a display of the portion or representation of the second segment on a computer display monitor.

27. A method as in claim 20, further comprising the step of identifying the subject matter content of a segment of the body of information, wherein the step of comparing further comprises the step of determining the similarity of the subject matter content of a segment to the subject matter content of a different segment, the predetermined criteria including a predefined degree of similarity with respect to which the relatedness of the compared segments is determined.

28. A method as in claim 27, wherein the step of determining the similarity of the subject matter of segments further comprises the step of performing a relevance feedback method.

29. A method as in claim 20, wherein the step of acquiring data further comprises the step of acquiring television broadcast signals.

30. A method as in claim 20, wherein the step of acquiring data further comprises the step of acquiring radio broadcast signals.

31. A method as in claim 20, wherein the step of acquiring data further comprises the step of acquiring computer-readable data files over a computer network from an information providing site that is part of that network.

32. A method as in claim 20, wherein the step of acquiring data further comprises the steps of:

40

acquiring television broadcast signals; and

acquiring computer-readable data files over a computer network from an information providing site that is part of that network.

33. A method as in claim 22, wherein:

the first segment is represented by data produced from the television broadcast signals; and

the second segment is represented by data from the computer-readable data files.

34. A method as in claim 20, further comprising the step of identifying an instruction from a user to begin displaying at least some of the body of information, wherein the display of a first segment is begun in response to the user instruction.

35. A method as in claim 20, wherein the first and second segments are displayed on physically separate display devices.

36. A method as in claim 20, wherein the steps of storing the acquired data, generating a display of a first segment of the body of information, and generating a display of a portion of, or a representation of, a second segment of the body of information are performed by devices interconnected to a conventional computer bus that enables the devices to communicate with each other such that the devices do not require wire communication over network communication lines to communicate with each other.

37. A method as in claim 20, wherein at least some of the acquired data is digital data, the step of acquiring data further comprising the step of acquiring digital data.

38. A method as in claim 20, wherein at least some of the acquired data is analog data, the step of acquiring data further comprising the step of acquiring analog data.

39. A method for categorizing according to subject matter an uncategorized segment of a body of information that includes a plurality of segments, each segment representing a defined set of information in the body of information, one or more segments of the body of information having previously been categorized by identifying each of the one or more segments with one or more subject matter categories, the method comprising the steps of:

determining the degree of similarity between the subject matter content of the uncategorized segment and the subject matter content of each of the previously categorized segments;

identifying one or more of the previously categorized segments as relevant to the uncategorized segment based upon the determined degrees of similarity of subject matter content between the uncategorized segment and the previously categorized segments; and

selecting one or more subject matter categories with which to identify the uncategorized segment based upon the subject matter categories used to identify the relevant previously categorized segments.

40. A method as in claim 39, wherein the step of determining the degree of similarity is accomplished using a relevance feedback method.

41. A method as in claim 39, wherein the step of identifying one or more of the previously categorized segments as relevant to the uncategorized segment further comprises the steps of:

identifying a plurality of the previously categorized segments that are the most similar to the uncategorized segment;

determining the degree of similarity between each of the plurality of previously categorized segments and each other of the plurality of previously categorized segments;

US 6,263,507 B1

41

for each pair of previously categorized segments of the plurality of previously categorized segments having greater than a predefined degree of similarity, eliminating one of the pair of previously categorized segments from the plurality of previously categorized segments, wherein the previously categorized segment or segments remaining after the step of eliminating are similar and distinct previously categorized segments; and

identifying one or more of the similar and distinct previously categorized segments as relevant previously categorized segments.

42. A method as in claim 39, wherein the step of selecting one or more subject matter categories further comprises selecting the most frequently occurring subject matter category or categories associated with the relevant previously categorized segments.

43. A method as in claim 39, wherein the uncategorized segment has been acquired from a first data source and the previously categorized segment or segments have been acquired from a second data source that is different than the first data source.

44. A method as in claim 43, wherein:

the data acquired from the first data source are television or radio broadcast signals; and

the data acquired from the second data source are computer-readable data files.

45. A method for determining whether a first set of information represented by a set of data of a first type is relevant to a second set of information represented by a set of data of a second type, the first and second sets of information being different from each other, the method comprising the steps of:

deriving a set of data of the second type from the set of data of the first type, the derived set of data of the second type also being representative of the first set of information;

determining the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information; and

determining whether the first set of information is relevant to the second set of information based upon the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information.

46. A method as in claim 45, wherein the first type of data is audiovisual data and the second type of data is text data.

47. A method as in claim 46, wherein the step of determining the degree of similarity is accomplished using a relevance feedback method.

48. A method as in claim 48, wherein a plurality of sets of information, each different from the other sets of the plurality of sets of information, are each represented by an associated set of data of the second type, the method enabling determination of which, if any, of the plurality of sets of information represented by a set of data of the second type are relevant to the first set of information represented by the set of data of the first type, the method further comprising the steps of:

determining the degree of similarity between each set of data of the second type representing one of the plurality of sets of information and the derived set of data of the second type representing the first set of information;

identifying which, if any, of the sets of data of the second type representing one of the plurality of sets of infor-

42

mation have greater than a predefined degree of similarity to the derived set of data of the second type representing the first set of information, the sets of data of the second type so identified being termed similar sets of data of the second type;

determining the degree of similarity between each similar set of data of the second type and each other similar set of data of the second type;

for each pair of similar sets of data of the second type having greater than a predefined degree of similarity, eliminating one of the pair of similar sets of data of the second type from the set of similar sets of data of the second type, wherein the set or sets of similar data of the second type remaining after the step of eliminating are similar and distinct sets of data of the second type; and

identifying the set or sets of information corresponding to one or more of the similar and distinct sets of data of the second type as relevant to the second set of information.

49. A method as in claim 48, wherein the step of identifying the relevant set or sets of information further comprises identifying no more than a predetermined number of relevant sets of information, the predetermined number of relevant sets of information corresponding to the sets of data of the second type having the greatest degree of similarity to the derived set of data of the second type.

50. A method as in claim 48, wherein the first type of data is analog data and the second type of data is digital data.

51. A method for identifying the boundaries of segments in a body of information, each segment comprising a contiguous related set of information in the body of information, wherein the body of information is represented by at least a set of text data and a set of video data, the method comprising the steps of:

performing a coarse partitioning method, the coarse partitioning method further comprising the steps of:

identifying time-stamped markers in the set of text data; and

determining approximate segment boundaries within the body of information as the times of occurrence of the time-stamp markers;

for each approximate segment boundary, specifying a range of time that includes the time of occurrence of the approximate segment boundary;

extracting subsets of video data from the set of video data that occur during the specified ranges of time;

performing a fine partitioning method to identify one or more breaks in the set of video data; and

selecting the best break that occurs in each subset of video data, the time of occurrence of the best break in each subset being designated as a boundary of a segment in the body of information.

52. A method as in claim 51, wherein the step of performing a fine partitioning method further comprises identifying the best breaks using a process that includes scene break identification.

53. A method as in claim 51, wherein the step of fine partitioning is performed on the entire set of video data to identify all of the breaks in the set of video data.

54. A method as in claim 51, wherein the step of fine partitioning is performed only on the subsets of video data to identify only breaks that occur in the subsets.

55. A method as in claim 51, wherein the best break of each subset is determined according to the criteria of the fine partitioning method used.

US 6,263,507 B1

43

56. A method as in claim 51, wherein the best break of each subset is the break occurring closest in time to the time of occurrence of the segment boundary in the text data that corresponds to that subset.

57. A method as in claim 51, wherein the body of information is represented by a set of text data, a set of audio data and a set of video data, the method further comprising the steps of:

ascertaining a synchronization of the audio data and the video data; and

determining the location of the segment boundaries in the set of audio data using the previously determined location of the segment boundaries in the set of video data and the synchronization of the audio data and video data.

58. A method for identifying the boundaries of segments in a body of information, each segment comprising a contiguous related set of information in the body of information, wherein the body of information is represented by a set of text data, a set of video data, and a set of audio data, the method comprising the steps of:

performing a coarse partitioning method, the coarse partitioning method further comprising the steps of:

identifying time-stamped markers in the set of text data; and

determining approximate segment boundaries within the body of information as the times of occurrence of the time-stamp markers;

for each approximate segment boundary, specifying a range of time that includes the time of occurrence of the approximate segment boundary;

extracting subsets of audio data from the set of audio data that occur during the specified ranges of time;

performing a fine partitioning method to identify one or more breaks in the set of audio data;

selecting the best break that occurs in each subset of audio data, the time of occurrence of the best break in each subset being designated as a boundary of a segment in the body of information;

ascertaining a synchronization of the audio data and the video data; and

determining the location of the segment boundaries in the set of video data using the previously determined location of the segment boundaries in the set of audio data and the synchronization of the audio data and video data.

59. A method as in claim 58, wherein the step of performing fine partitioning further comprises identifying the best breaks using a process that includes pause recognition.

60. A method as in claim 58, wherein the step of performing fine partitioning further comprises identifying the best breaks using a process that includes voice recognition.

61. A method as in claim 58, wherein the step of performing fine partitioning further comprises identifying the best breaks using a process that includes word recognition.

62. A method as in claim 58, wherein the step of performing fine partitioning further comprises identifying the best breaks using a process that includes music recognition.

63. A computer readable medium encoded with one or more computer programs for enabling acquisition and review of a body of information, wherein the body of information includes a plurality of segments, each segment representing a defined set of information in the body of information, comprising:

instructions for acquiring data representing the body of information;

44

instructions for storing the acquired data;

instructions for generating a display of a first segment of the body of information from data that is part of the stored data;

instructions for comparing data representing a segment of the body of information to data representing a different segment of the body of information to determine whether, according to one or more predetermined criteria, the compared segments are related; and

instructions for generating a display of a portion of, or a representation of, a second segment of the body of information from data that is part of the stored data, wherein the display of the portion or representation of the second segment is generated in response to the display of a first segment to which the second segment is related.

64. A computer readable medium as in claim 63, further comprising instructions for causing the display of the portion or representation of the second segment to occur substantially coextensive in time with the display of the related first segment.

65. A computer readable medium as in claim 63, wherein: the instructions for acquiring data representing the body of information further comprise instructions for acquiring audiovisual data representing at least a portion of the body of information, wherein the first and second segments are represented by audiovisual data; and

the instructions for generating a display of a first segment of the body of information further comprise instruction for generating an audiovisual display of the first segment.

66. A computer readable medium as in claim 65, further comprising instructions for identifying the selection of a second segment for which a portion or representation is being displayed, wherein selection of such second segment causes an audiovisual display of the selected second segment to be produced.

67. A computer readable medium as in claim 63, wherein: the instructions for acquiring data representing the body of information further comprise instructions for acquiring audiovisual data representing at least a portion of the body of information;

the instructions for generating a display of a first segment of the body of information further comprise instructions for generating an audiovisual display of the first segment; and

the instructions for generating a display of a portion of, or a representation of, a second segment of the body of information further comprise instructions for generating a text display of the portion or representation of the second segment.

68. A computer readable medium as in claim 63, wherein: the instructions for generating a display of a first segment of the body of information further comprise instructions for generating a display of the first segment on an analog display device; and

the instructions for generating a display of a portion of, or a representation of, a second segment of the body of information further comprise instructions for generating a display of the portion or representation of the second segment on a digital display device.

69. A computer readable medium as in claim 63, wherein: the instructions for generating a display of the first segment on an analog display device further comprise instructions for generating a display of the first segment on a television; and

US 6,263,507 B1

45

the instructions for generating a display of the portion or representation of the second segment on a digital display device further comprise instructions for generating a display of the portion or representation of the second segment on a computer display monitor.

70. A computer readable medium as in claim 63, further comprising instructions for identifying the subject matter content of a segment of the body of information, wherein the instructions for comparing further comprise instructions for determining the similarity of the subject matter content of a segment to the subject matter content of a different segment, the predetermined criteria including a predefined degree of similarity with respect to which the relatedness of the compared segments is determined.

71. A computer readable medium as in claim 70, wherein the instructions for determining the similarity of the subject matter of segments further comprise instructions for performing a relevance feedback method.

72. A computer readable medium as in claim 63, wherein the instructions for acquiring data further comprise instructions for acquiring television broadcast signals.

73. A computer readable medium as in claim 63, wherein the instructions for acquiring data further comprise instructions for acquiring radio broadcast signals.

74. A computer readable medium as in claim 63, wherein the instructions for acquiring data further comprise instructions for acquiring computer-readable data files over a computer network from an information providing site that is part of that network.

75. A computer readable medium as in claim 63, wherein the instructions for acquiring data further comprise:

instructions for acquiring television broadcast signals; and
instructions for acquiring computer-readable data files over a computer network from an information providing site that is part of that network.

76. A computer readable medium as in claim 75, wherein: the first segment is represented by data produced from the television broadcast signals; and
the second segment is represented by data from the computer-readable data files.

77. A computer readable medium as in claim 63, further comprising instructions for identifying an instruction from a user to begin displaying at least some of the body of information, wherein the display of a first segment is begun in response to the user instruction.

78. A computer readable medium as in claim 63, wherein the first and second segments are displayed on physically separate display devices.

79. A computer readable medium as in claim 63, wherein the instructions for storing the acquired data, generating a display of a first segment of the body of information, and generating a display of a portion of, or a representation of, a second segment of the body of information are executed by devices interconnected to a conventional computer bus that enables the devices to communicate with each other such that the devices do not require wire communication over network communication lines to communicate with each other.

80. A computer readable medium as in claim 63, wherein at least some of the acquired data is digital data, the instructions for acquiring data further comprising instructions for acquiring digital data.

81. A computer readable medium as in claim 63, wherein at least some of the acquired data is analog data, the instructions for acquiring data further comprising instructions for acquiring analog data.

46

82. A computer readable medium encoded with one or more computer programs for enabling categorization according to subject matter of an uncategorized segment of a body of information that includes a plurality of segments, each segment representing a defined set of information in the body of information, one or more segments having previously been categorized by identifying each of the one or more segments with one or more subject matter categories, comprising:

instructions for determining the degree of similarity between the subject matter content of the uncategorized segment and the subject matter content of each of the previously categorized segments;

instructions for identifying one or more of the previously categorized segments as relevant to the uncategorized segment based upon the determined degrees of similarity of subject matter content between the uncategorized segment and the previously categorized segments; and

instructions for selecting one or more subject matter categories with which to identify the uncategorized segment based upon the subject matter categories used to identify the relevant previously categorized segments.

83. A computer readable medium as in claim 82, wherein the instructions for determining the degree of similarity further comprise instructions for performing a relevance feedback method.

84. A computer readable medium as in claim 82, wherein the instructions for identifying one or more of the previously categorized segments as relevant to the uncategorized segment further comprise:

instructions for identifying a plurality of the previously categorized segments that are the most similar to the uncategorized segment;

instructions for determining the degree of similarity between each of the plurality of previously categorized segments and each other of the plurality of previously categorized segments;

instructions for eliminating, for each pair of previously categorized segments of the plurality of previously categorized segments having greater than a predefined degree of similarity, one of the pair of previously categorized segments from the plurality of previously categorized segments, wherein the remaining previously categorized segment or segments are similar and distinct previously categorized segments; and

instructions for identifying one or more of the similar and distinct previously categorized segments as relevant previously categorized segments.

85. A computer readable medium as in claim 82, wherein the instructions for selecting one or more subject matter categories further comprise instructions for selecting the most frequently occurring subject matter category or categories associated with the relevant previously categorized segments.

86. A computer readable medium as in claim 82, wherein the uncategorized segment has been acquired from a first data source and the previously categorized segment or segments have been acquired from a second data source that is different than the first data source.

87. A computer readable medium as in claim 86, wherein: the data acquired from the first data source are television or radio broadcast signals; and
the data acquired from the second data source are computer-readable data files.

US 6,263,307 B1

47

88. A computer readable medium encoded with one or more computer programs for enabling determination of whether a first set of information represented by a set of data of a first type is relevant to a second set of information represented by a set of data of a second type, the first and second sets of information being different from each other, comprising:

instructions for deriving a set of data of the second type from the set of data of the first type, the derived set of data of the second type also being representative of the first set of information;

instructions for determining the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information; and

instructions for determining whether the first set of information is relevant to the second set of information based upon the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information.

89. A computer readable medium as in claim 88, wherein the first type of data is audiovisual data and the second type of data is text data.

90. A computer readable medium as in claim 89, wherein the instructions for determining the degree of similarity further comprise instructions for performing a relevance feedback method.

91. A computer readable medium as in claim 88, wherein a plurality of sets of information, each different from the other sets of the plurality of sets of information, are each represented by an associated set of data of the second type, the one or more computer programs enabling determination of which, if any, of the plurality of sets of information represented by a set of data of the second type are relevant to the first set of information represented by the set of data of the first type, the one or more computer programs further comprising:

instructions for determining the degree of similarity between each set of data of the second type representing one of the plurality of sets of information and the derived set of data of the second type representing the first set of information;

instructions for identifying which, if any, of the sets of data of the second type representing one of the plurality of sets of information have greater than a predefined degree of similarity to the derived set of data of the second type representing the first set of information, the sets of data of the second type so identified being termed similar sets of data of the second type;

instructions for determining the degree of similarity between each similar set of data of the second type and each other similar set of data of the second type;

instructions for eliminating, for each pair of similar sets of data of the second type having greater than a predefined degree of similarity, one of the pair of similar sets of data of the second type from the set of similar sets of data of the second type, wherein the remaining set or sets of similar data of the second type are similar and distinct sets of data of the second type; and

instructions for identifying the set or sets of information corresponding to one or more of the similar and distinct sets of data of the second type as relevant to the second set of information.

92. A computer readable medium as in claim 91, wherein the instructions for identifying the relevant set or sets of

48

information further comprise instructions for identifying no more than a predetermined number of relevant sets of information, the predetermined number of relevant sets of information corresponding to the sets of data of the second type having the greatest degree of similarity to the derived set of data of the second type.

93. A computer readable medium as in claim 88, wherein the first type of data is analog data and the second type of data is digital data.

94. A computer readable medium encoded with one or more computer programs for enabling identification of the boundaries of segments in a body of information, each segment comprising a contiguous related set of information in the body of information, wherein the body of information is represented by at least a set of text data and a set of video data, comprising:

instructions for performing a coarse partitioning method, the coarse partitioning instructions further comprising: instructions for identifying time-stamped markers in the set of text data; and

instructions for determining approximate segment boundaries within the body of information as the times of occurrence of the time-stamp markers;

instructions for specifying, for each approximate segment boundary, a range of time that includes the time of occurrence of the approximate segment boundary;

instructions for extracting subsets of video data from the set of video data that occur during the specified ranges of time;

instructions for performing a fine partitioning method to identify one or more breaks in the set of video data; and instructions for selecting the best break that occurs in each subset of video data, the time of occurrence of the best break in each subset being designated as a boundary of a segment in the body of information.

95. A computer readable medium as in claim 94, wherein the instructions for performing a fine partitioning method further comprise instructions for identifying the best breaks using a process that includes scene break identification.

96. A computer readable medium as in claim 94, wherein the fine partitioning method is performed on the entire set of video data to identify all of the breaks in the set of video data.

97. A computer readable medium as in claim 94, wherein the fine partitioning method is performed only on the subsets of video data to identify only breaks that occur in the subsets.

98. A computer readable medium as in claim 94, wherein the best break of each subset is determined according to the criteria of the fine partitioning method used.

99. A computer readable medium as in claim 94, wherein the best break of each subset is the break occurring closest in time to the time of occurrence of the segment boundary in the text data that corresponds to that subset.

100. A computer readable medium as in claim 94, wherein the body of information is represented by a set of text data, a set of audio data and a set of video data, the one or more computer programs further comprising:

instructions for ascertaining a synchronization of the audio data and the video data; and

instructions for determining the location of the segment boundaries in the set of audio data using the previously determined location of the segment boundaries in the set of video data and the synchronization of the audio data and video data.

101. A system for categorizing according to subject matter an uncategorized segment of a body of information that

US 6,263,507 B1

49

includes a plurality of segments, each segment representing a defined set of information in the body of information, one or more segments of the body of information having previously been categorized by identifying each of the one or more segments with one or more subject matter categories, the system comprising:

means for determining the degree of similarity between the subject matter content of the uncategorized segment and the subject matter content of each of the previously categorized segments;

means for identifying one or more of the previously categorized segments as relevant to the uncategorized segment based upon the determined degrees of similarity of subject matter content between the uncategorized segment and the previously categorized segments; and

means for selecting one or more subject matter categories with which to identify the uncategorized segment based upon the subject matter categories used to identify the relevant previously categorized segments.

103. A system as in claim 101, wherein the means for determining the degree of similarity further comprises means for performing a relevance feedback method.

103. A system as in claim 101, wherein the means for identifying one or more of the previously categorized segments as relevant to the uncategorized segment further comprises:

means for identifying a plurality of the previously categorized segments that are the most similar to the uncategorized segment;

means for determining the degree of similarity between each of the plurality of previously categorized segments and each other of the plurality of previously categorized segments;

means for eliminating, for each pair of previously categorized segments of the plurality of previously categorized segments having greater than a predefined degree of similarity, one of the pair of previously categorized segments from the plurality of previously categorized segments, wherein the remaining previously categorized segment or segments are similar and distinct previously categorized segments; and

means for identifying one or more of the similar and distinct previously categorized segments as relevant previously categorized segments.

104. A system as in claim 101, wherein the means for selecting one or more subject matter categories further comprises means for selecting the most frequently occurring subject matter category or categories associated with the relevant previously categorized segments.

105. A system as in claim 101, wherein the uncategorized segment has been acquired from a first data source and the previously categorized segment or segments have been acquired from a second data source that is different than the first data source.

106. A system as in claim 105, wherein:

the data acquired from the first data source are television or radio broadcast signals; and

the data acquired from the second data source are computer-readable data files.

107. A system for determining whether a first set of information represented by a set of data of a first type is relevant to a second set of information represented by a set of data of a second type, the first and second sets of information being different from each other, the system comprising:

50

means for deriving a set of data of the second type from the set of data of the first type, the derived set of data of the second type also being representative of the first set of information;

means for determining the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information; and

means for determining whether the first set of information is relevant to the second set of information based upon the degree of similarity between the set of data of the second type representing the second set of information and the derived set of data of the second type representing the first set of information.

108. A system as in claim 107, wherein the first type of data is audiovisual data and the second type of data is text data.

109. A system as in claim 108, wherein the means for determining the degree of similarity further comprises means for performing a relevance feedback method.

110. A system as in claim 107, wherein a plurality of sets of information, each different from the other sets of the plurality of sets of information, are each represented by an associated set of data of the second type, the system enabling determination of which, if any, of the plurality of sets of information represented by a set of data of the second type are relevant to the first set of information represented by the set of data of the first type, the system further comprising:

means for determining the degree of similarity between each set of data of the second type representing one of the plurality of sets of information and the derived set of data of the second type representing the first set of information;

means for identifying which, if any, of the sets of data of the second type representing one of the plurality of sets of information have greater than a predefined degree of similarity to the derived set of data of the second type representing the first set of information, the sets of data of the second type so identified being termed similar sets of data of the second type;

means for determining the degree of similarity between each similar set of data of the second type and each other similar set of data of the second type;

means for eliminating, for each pair of similar sets of data of the second type having greater than a predefined degree of similarity, one of the pair of similar sets of data of the second type from the set of similar sets of data of the second type, wherein the remaining set or sets of similar data of the second type are similar and distinct sets of data of the second type; and

means for identifying the set or sets of information corresponding to one or more of the similar and distinct sets of data of the second type as relevant to the second set of information.

111. A system as in claim 110, wherein the means for identifying the relevant set or sets of information further comprises means for identifying no more than a predetermined number of relevant sets of information, the predetermined number of relevant sets of information corresponding to the sets of data of the second type having the greatest degree of similarity to the derived set of data of the second type.

112. A system as in claim 107, wherein the first type of data is analog data and the second type of data is digital data.

113. A computer readable medium encoded with one or more computer programs for identifying the boundaries of

US 6,263,507 B1

51

segments in a body of information, each segment comprising a contiguous related set of information in the body of information, wherein the body of information is represented by a set of text data, a set of video data, and a set of audio data, comprising:

instructions for performing a coarse partitioning method, the instructions for performing a coarse partitioning method further comprising:

instructions for identifying time-stamped markers in the set of text data; and

instructions for determining approximate segment boundaries within the body of information as the times of occurrence of the time-stamp markers;

instructions for specifying, for each approximate segment boundary, a range of time that includes the time of occurrence of the approximate segment boundary;

instructions for extracting subsets of audio data from the set of audio data that occur during the specified ranges of time;

instructions for performing a fine partitioning method to identify one or more breaks in the set of audio data;

instructions for selecting the best break that occurs in each subset of audio data, the time of occurrence of the best break in each subset being designated as a boundary of a segment in the body of information;

instructions for ascertaining a synchronization of the audio data and the video data; and

instructions for determining the location of the segment boundaries in the set of video data using the previously determined location of the segment boundaries in the set of audio data and the synchronization of the audio data and video data.

114. A computer readable medium as in claim 113, wherein the instructions for performing fine partitioning further comprise instructions for identifying the best breaks using a process that includes pause recognition.

115. A computer readable medium as in claim 113, wherein the instructions for performing fine partitioning further comprise instructions for identifying the best breaks using a process that includes voice recognition.

116. A computer readable medium as in claim 113, wherein the instructions for performing fine partitioning further comprise instructions for identifying the best breaks using a process that includes word recognition.

117. A computer readable medium as in claim 113, wherein the instructions for performing fine partitioning further comprise instructions for identifying the best breaks using a process that includes music recognition.

118. A system for identifying the boundaries of segments in a body of information, each segment comprising a contiguous related set of information in the body of information, wherein the body of information is represented by at least a set of text data and a set of video data, the system comprising:

means for performing a coarse partitioning method, the means for performing a coarse partitioning method further comprising:

means for identifying time-stamped markers in the set of text data; and

means for determining approximate segment boundaries within the body of information as the times of occurrence of the time-stamp markers;

means for specifying, for each approximate segment boundary, a range of time that includes the time of occurrence of the approximate segment boundary;

means for extracting subsets of video data from the set of video data that occur during the specified ranges of time;

52

means for performing a fine partitioning method to identify one or more breaks in the set of video data; and
means for selecting the best break that occurs in each subset of video data, the time of occurrence of the best break in each subset being designated as a boundary of a segment in the body of information.

119. A system as in claim 118, wherein the means for performing a fine partitioning method further comprises means for identifying the best breaks using a process that includes scene break identification.

120. A system as in claim 118, wherein the fine partitioning method is performed on the entire set of video data to identify all of the breaks in the set of video data.

121. A system as in claim 118, wherein the fine partitioning method is performed only on the subsets of video data to identify only breaks that occur in the subsets.

122. A system as in claim 118, wherein the best break of each subset is determined according to the criteria of the fine partitioning method used.

123. A system as in claim 118, wherein the best break of each subset is the break occurring closest in time to the time of occurrence of the segment boundary in the text data that corresponds to that subset.

124. A system as in claim 118, wherein the body of information is represented by a set of text data, a set of audio data and a set of video data, the system further comprising:
means for ascertaining a synchronization of the audio data and the video data; and

means for determining the location of the segment boundaries in the set of audio data using the previously determined location of the segment boundaries in the set of video data and the synchronization of the audio data and video data.

125. A system for identifying the boundaries of segments in a body of information, each segment comprising a contiguous related set of information in the body of information, wherein the body of information is represented by a set of text data, a set of video data, and a set of audio data, the system comprising:

means for performing a coarse partitioning method, the means for performing a coarse partitioning method further comprising:

means for identifying time-stamped markers in the set of text data; and

means for determining approximate segment boundaries within the body of information as the times of occurrence of the time-stamp markers;

means for specifying, for each approximate segment boundary, a range of time that includes the time of occurrence of the approximate segment boundary;

means for extracting subsets of audio data from the set of audio data that occur during the specified ranges of time;

means for performing a fine partitioning method to identify one or more breaks in the set of audio data;

means for selecting the best break that occurs in each subset of audio data, the time of occurrence of the best break in each subset being designated as a boundary of a segment in the body of information;

means for ascertaining a synchronization of the audio data and the video data; and

means for determining the location of the segment boundaries in the set of video data using the previously determined location of the segment boundaries in the set of audio data and the synchronization of the audio data and video data.

US 6,263,507 B1

53

126. A system as in claim 125, wherein the means for performing a fine partitioning method further comprises means for identifying the best breaks using a process that includes pause recognition.

127. A system as in claim 125, wherein the means for performing a fine partitioning method further comprises means for identifying the best breaks using a process that includes voice recognition.

128. A system as in claim 125, wherein the means for performing a fine partitioning method further comprises

54

means for identifying the best breaks using a process that includes word recognition.

129. A system as in claim 125, wherein the means for performing a fine partitioning method further comprises means for identifying the best breaks using a process that includes music recognition.

* * * * *

EXHIBIT 3


US006034652A

United States Patent [19]**Freiberger et al.**[11] **Patent Number:** **6,034,652**[45] **Date of Patent:** ***Mar. 7, 2000**[54] **ATTENTION MANAGER FOR OCCUPYING THE PERIPHERAL ATTENTION OF A PERSON IN THE VICINITY OF A DISPLAY DEVICE****FOREIGN PATENT DOCUMENTS**WO 93/19427 9/1993 WIPO.
WO 96/30864 10/1996 WIPO.[75] **Inventors:** Paul A. Freiberger, San Mateo, Calif.;
Golan Levin, Staten Island, N.Y.;
David R. Reed, Atherton, Calif.; Marc
E. Davis, San Francisco, Calif.; Neal A.
Bhadhanakar, Philippe P. Parnet,
both of Palo Alto, Calif.; Todd A.
Agulnick, San Francisco, Calif.; Sally
N. Rosenthal, Palo Alto, Calif.; Giles
N. Goodfriend, Los Angeles, Calif.**OTHER PUBLICATIONS**Gomes, Lee, "Upstart's Internet 'TV' Has Microsoft Tuned In", *Wall Street Journal*, Aug. 1996.Joan E. Rigdon, "Screen Savers Go Beyond Fish, Flying Toasters," *Wall Street Journal*, Feb. 13, 1996.Staff Reporter, "PointCast Inc. Is Testing New Screen-Saver Product," *Wall Street Journal*, May 1996.[73] **Assignee:** Interval Research Corporation, Palo Alto, Calif.[*] **Notice:** This patent issued on a continued prosecution application filed under 37 CFR 1.53(d), and is subject to the twenty year patent term provisions of 35 U.S.C. 154(a)(2).**Primary Examiner—Jeffery Brier**
Attorney, Agent, or Firm—David R. Graham[57] **ABSTRACT**

An attention manager presents information to a person in the vicinity of a display device in a manner that engages at least the peripheral attention of the person. The information is embodied by one or more sets of content data (e.g., video or audio data). Each set of content data is formulated by a content provider and made available for use by content display systems. Upon appropriate activation, each content display system displays images corresponding to the sets of content data in accordance with predetermined scheduling information. The attention manager makes use of "unused capacity" of the display device and the person's attention, providing information to the person that the person might not otherwise expend adequate energy to obtain. The attention manager also affords an opportunity to content providers to disseminate their information to people that are interested in receiving such information, enabling the content providers to provide better directed information dissemination, as well as providing access to the previously unused attention capacity of those interested people.

[21] **Appl. No.:** 08/630,641[22] **Filed:** **Mar. 23, 1996**[51] **Int. Cl.⁷** G09G 5/12; G09G 5/14;
G06F 15/16[52] **U.S. Cl.** 348/2; 345/115; 709/218[58] **Field of Search** 345/2, 115, 116,
345/302, 326; 707/10, 104, 501; 395/200.47,
200.48, 200.49; 709/217, 218, 219[56] **References Cited****U.S. PATENT DOCUMENTS**

4,845,658	7/1989	Gifford	395/200.47
5,105,184	4/1992	Pirani et al.	345/115
5,305,195	4/1994	Murphy	705/1
5,347,632	9/1994	Flepp et al.	395/200.32
5,573,643	11/1996	Judson	395/200.48
5,740,549	4/1998	Reilly et al.	705/14
5,768,528	6/1998	Stamm	395/200.61

18 Claims, 8 Drawing Sheets

U.S. Patent

Mar. 7, 2000

Sheet 1 of 8

6,034,652

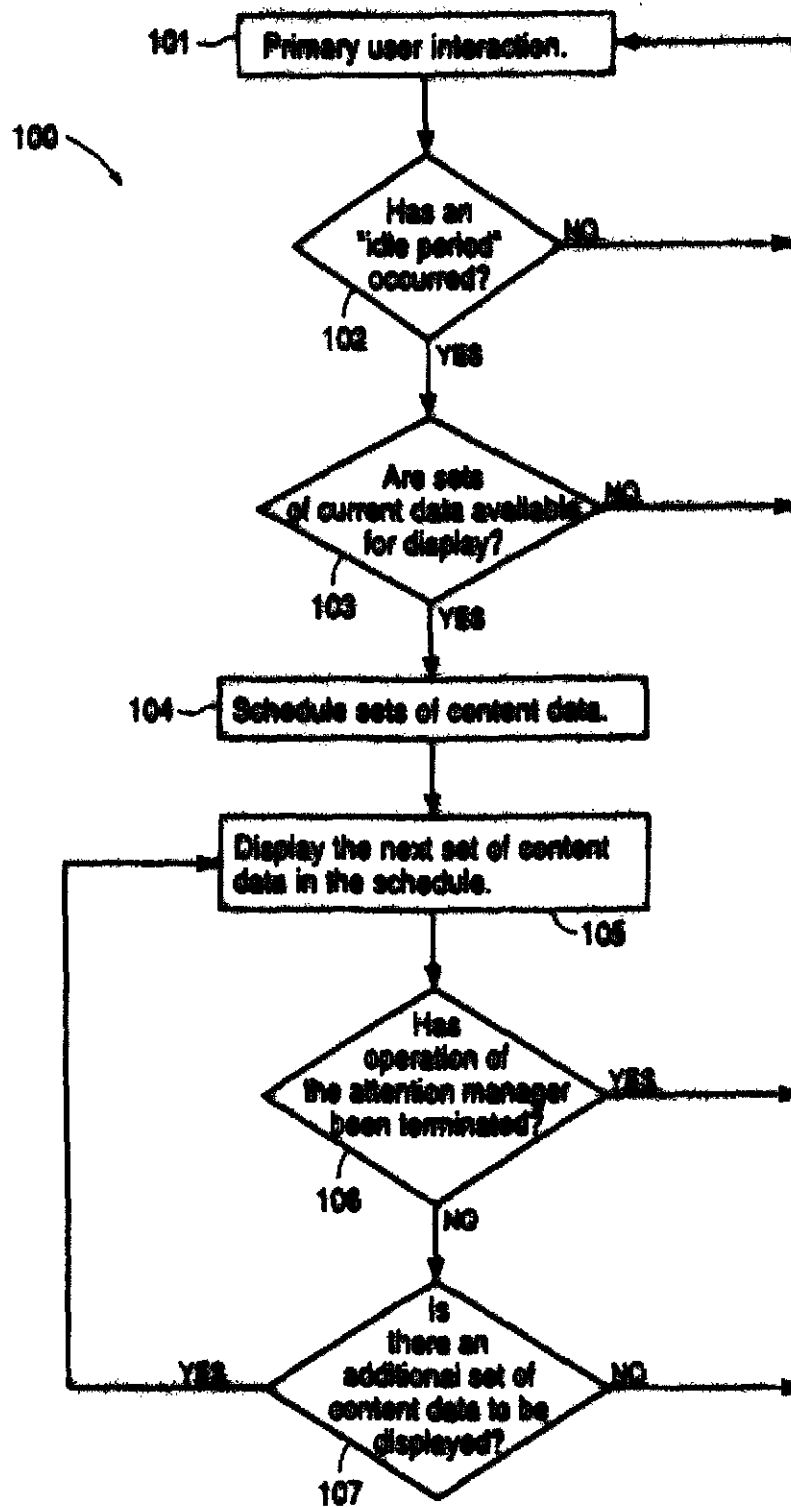


FIG. 1

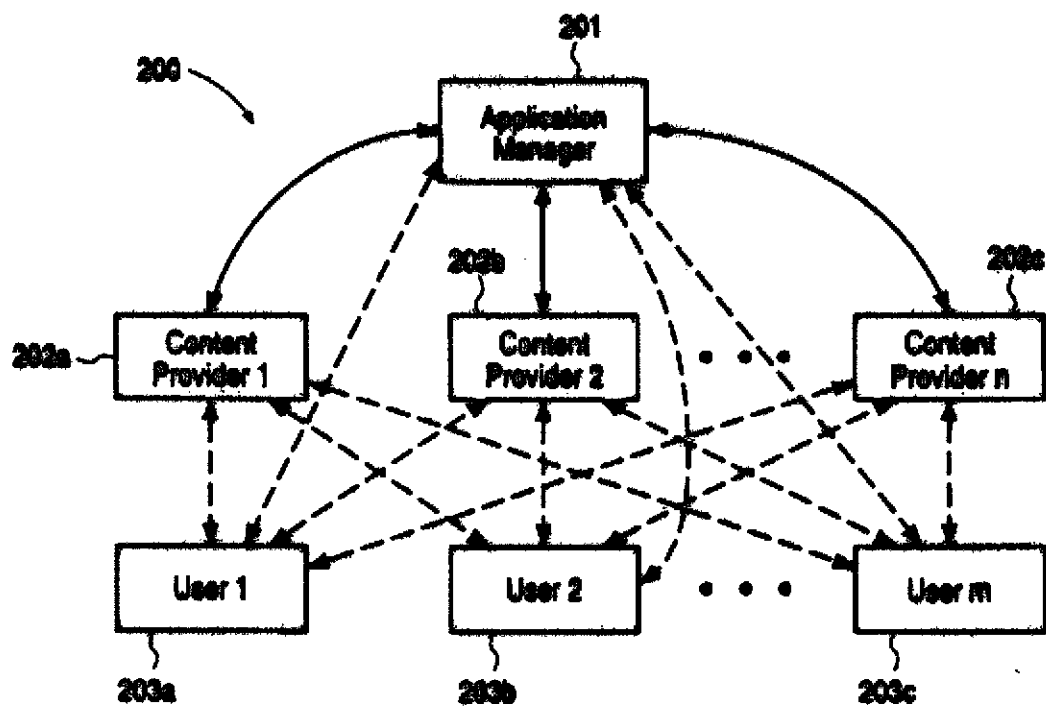


FIG. 2

U.S. Patent

Mar. 7, 2000

Sheet 3 of 8

6,034,652

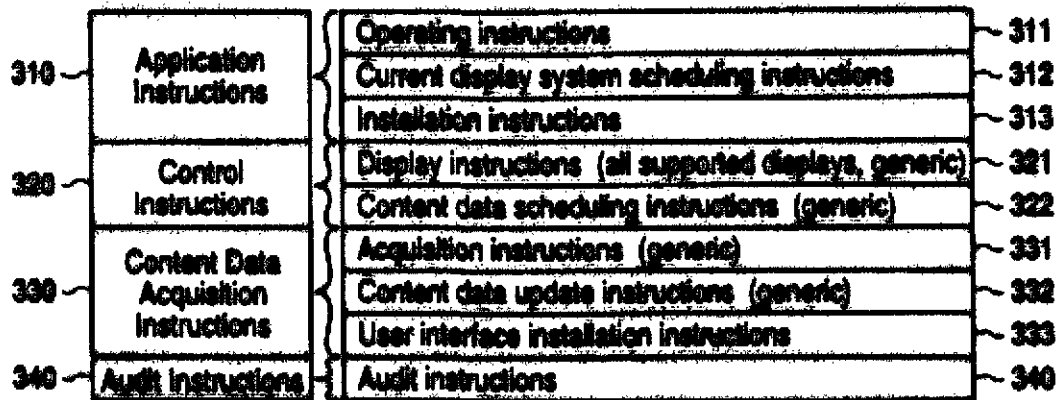


FIG. 3A

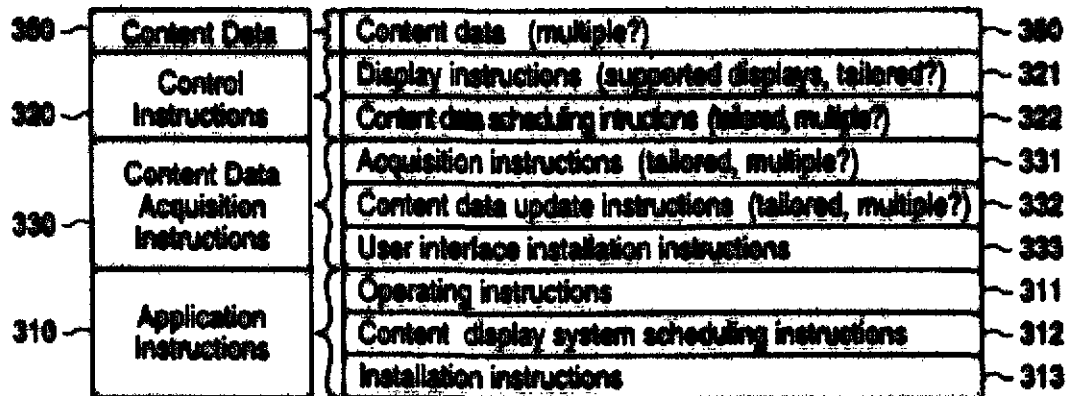


FIG. 3B

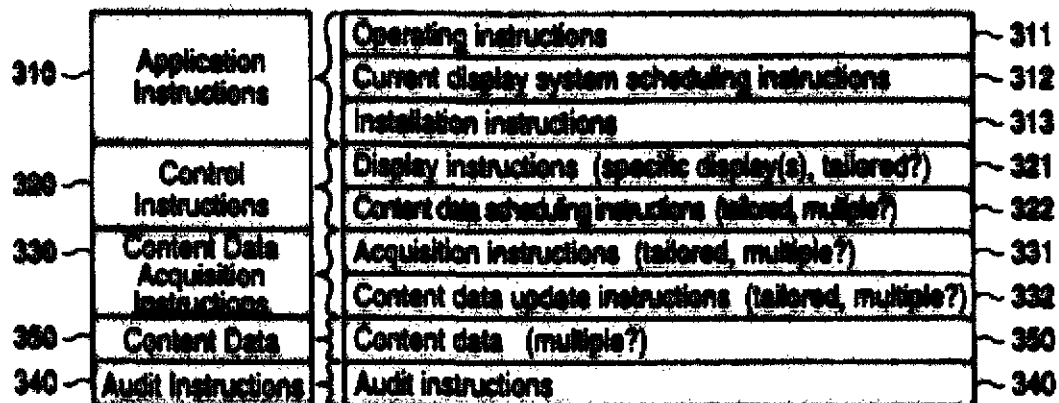


FIG. 3C

U.S. Patent

Mar. 7, 2000

Sheet 4 of 8

6,034,652

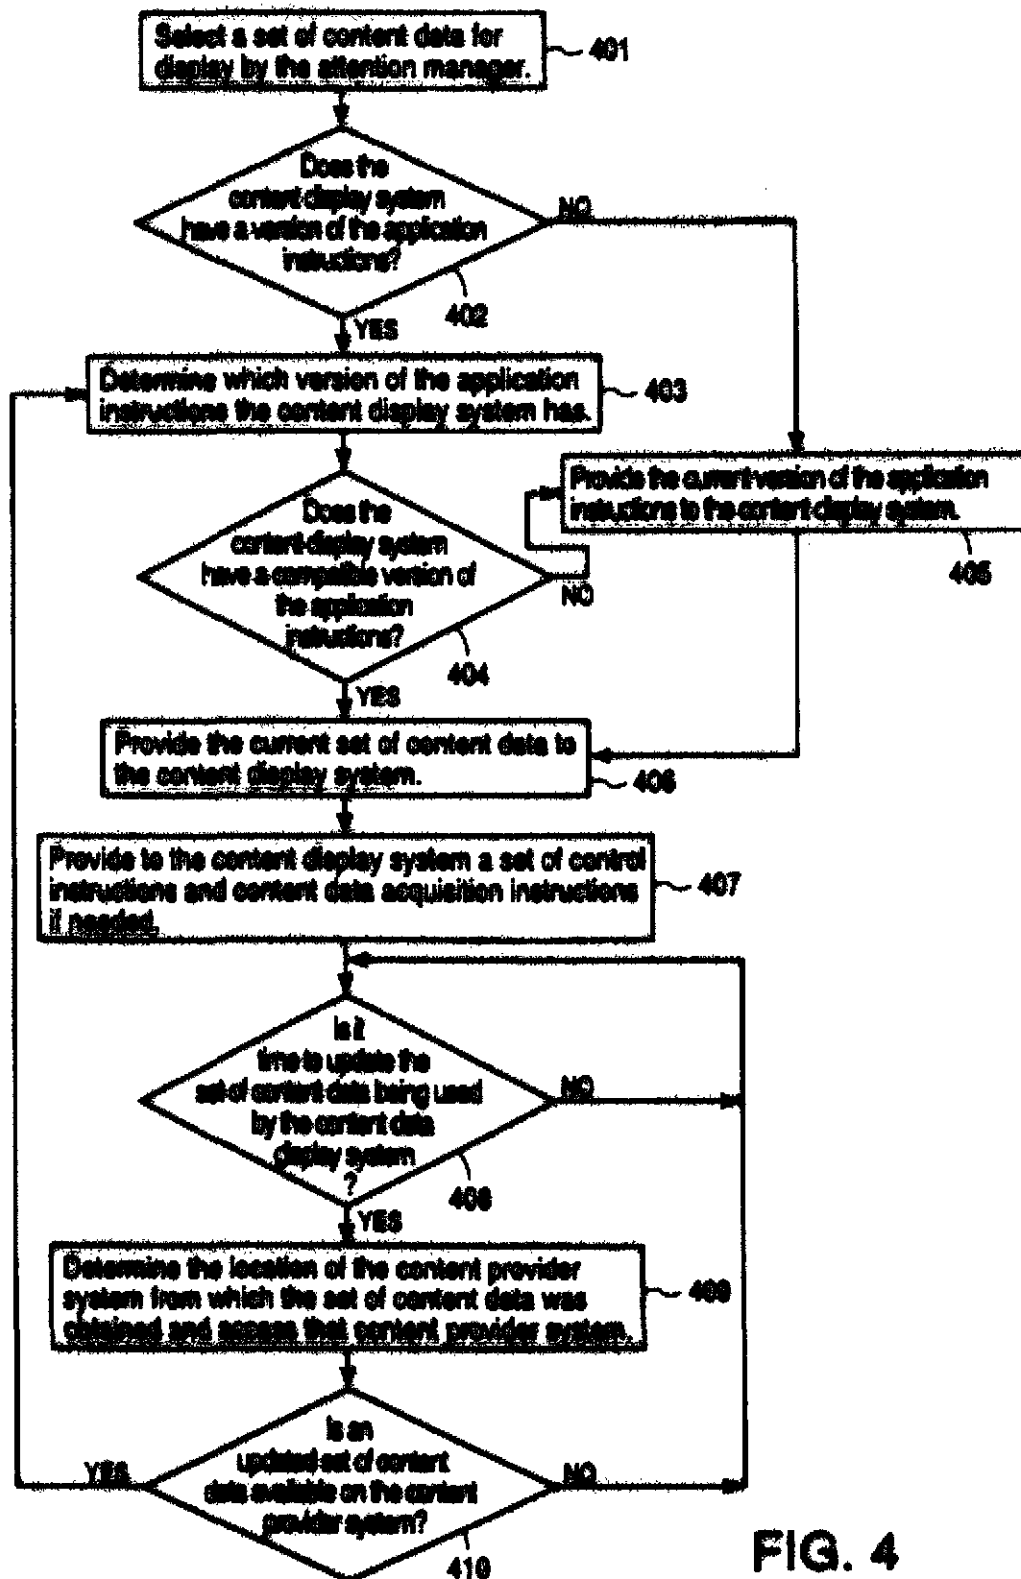


FIG. 4

FIG. 5A

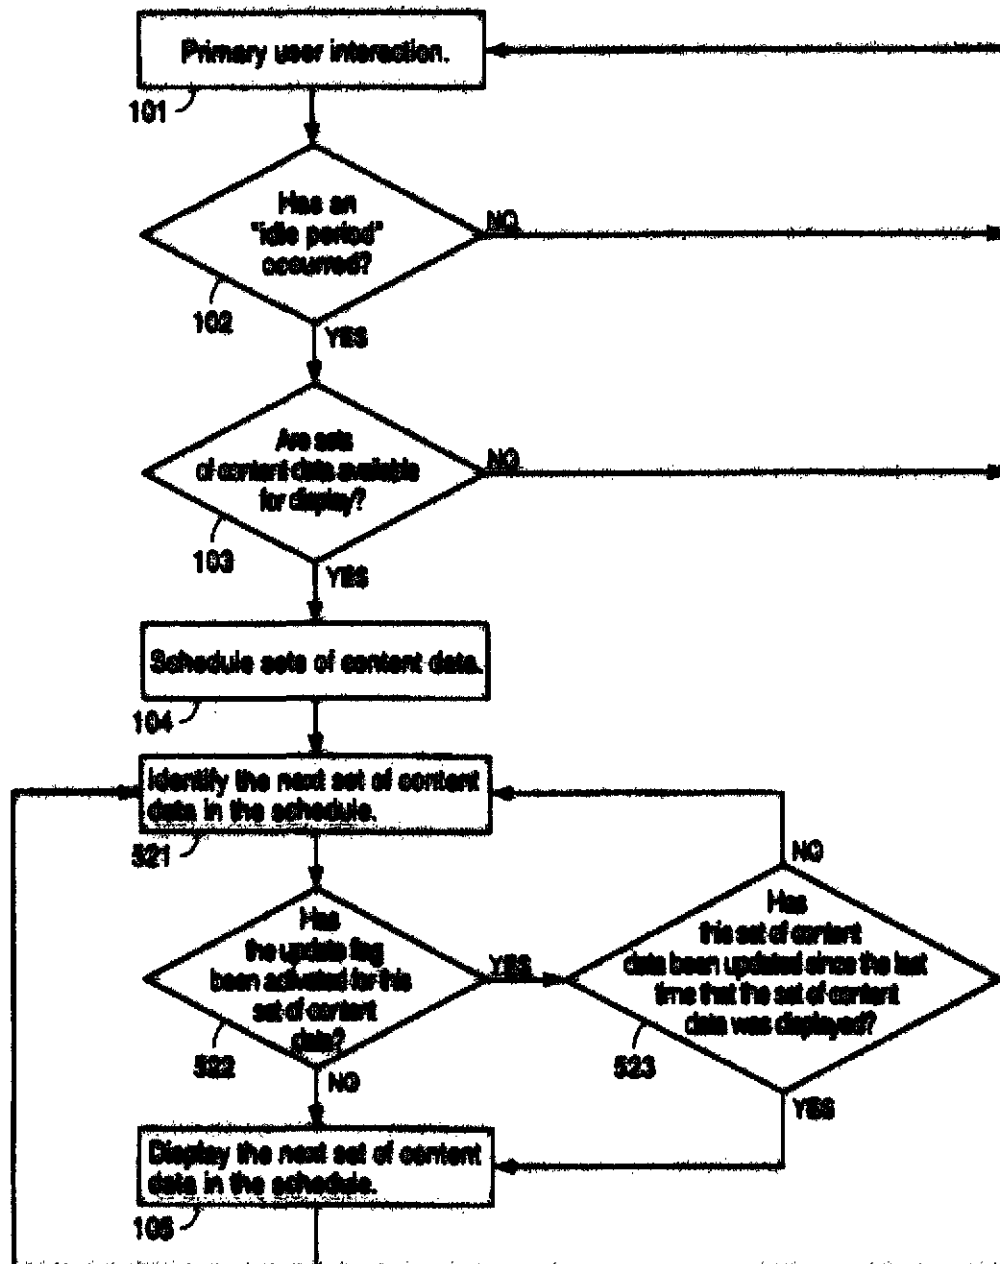
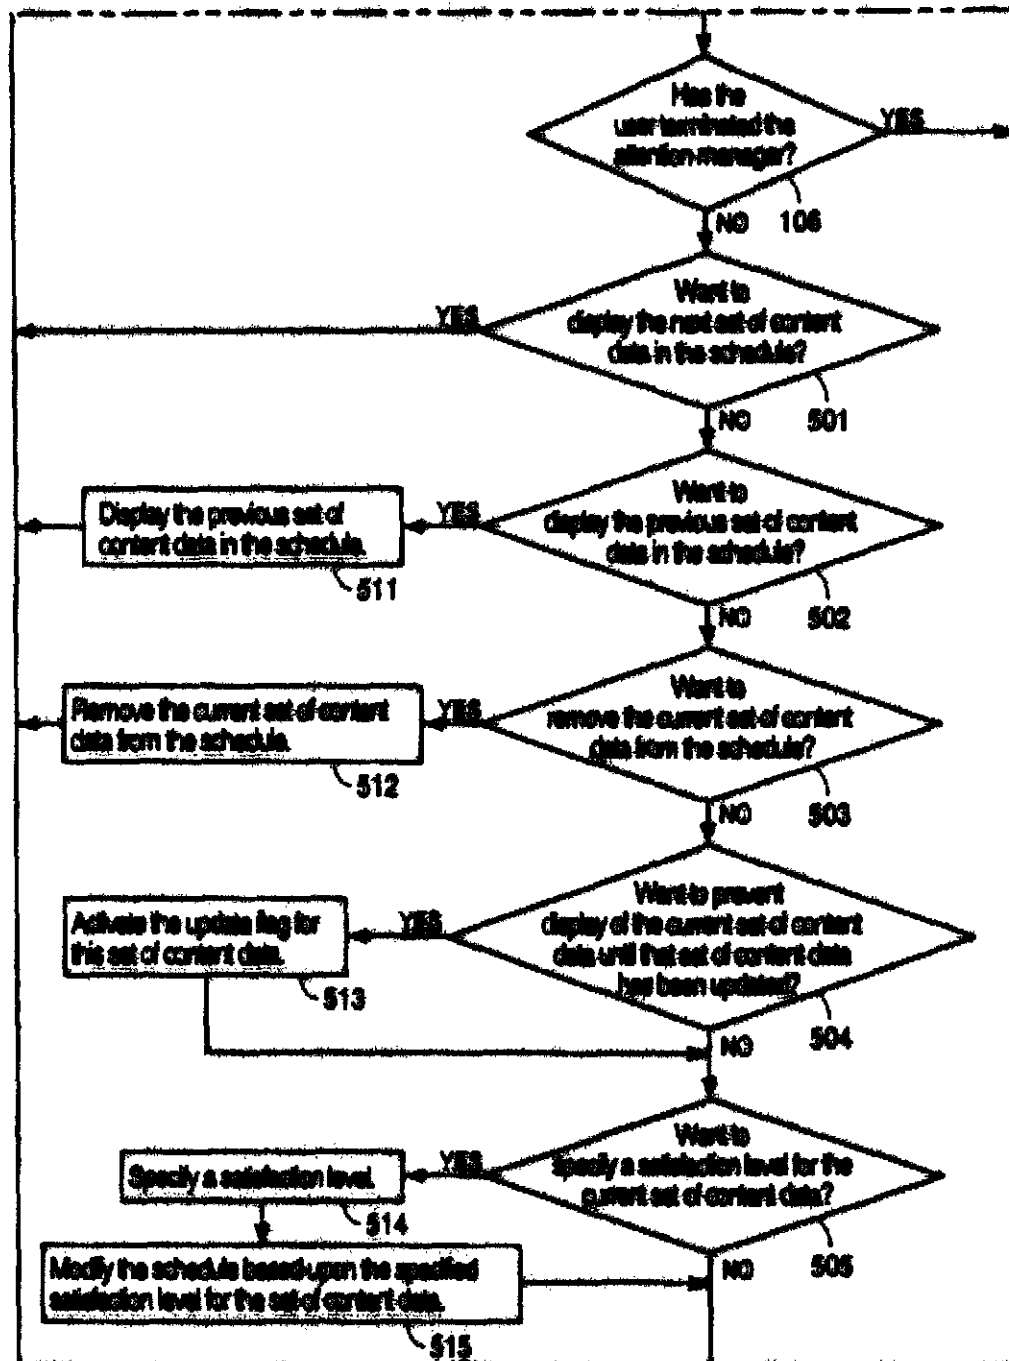


FIG. 5B



U.S. Patent

Mar. 7, 2000

Sheet 7 of 8

6,034,652

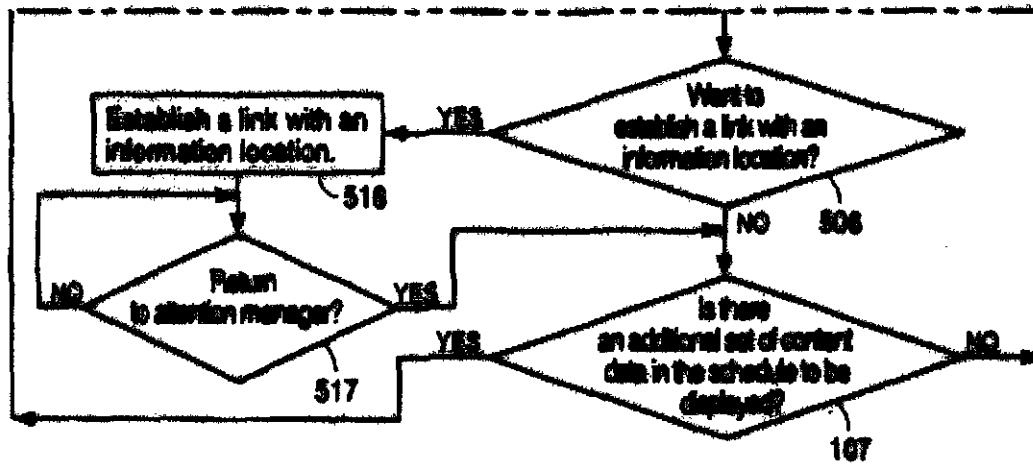


FIG. 5C



FIG. 5

U.S. Patent

Mar. 7, 2000

Sheet 8 of 8

6,034,652

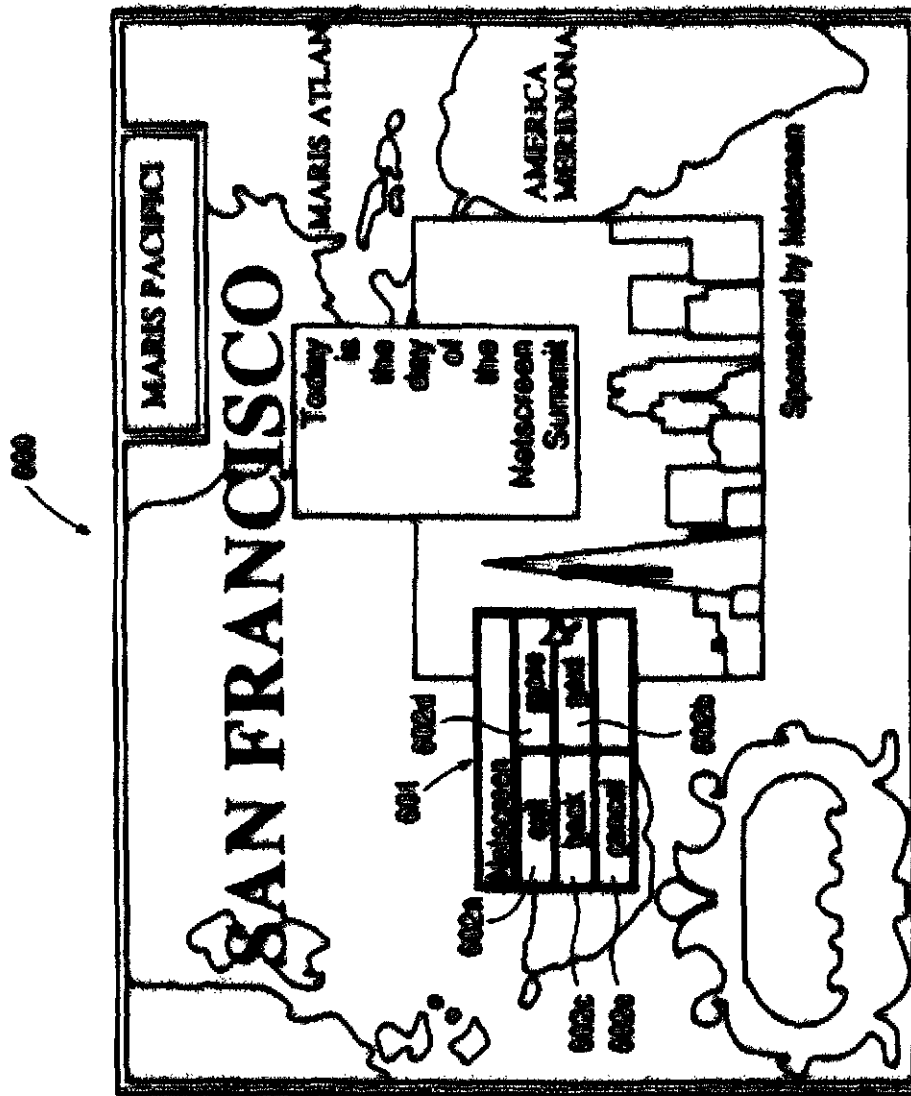


FIG. 6

6,034,652

1

ATTENTION MANAGER FOR OCCUPYING THE PERIPHERAL ATTENTION OF A PERSON IN THE VICINITY OF A DISPLAY DEVICE

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to the engagement of the peripheral attention of a person in the vicinity of a display device such as the display monitor of a computer.

2. Related Art

Information providers of all sorts have an interest in presenting their information to information consumers and, in particular, to information consumers who may, or do, have an interest in the particular information provided by the particular information provider. At the same time, information consumers have an interest in accessing a wide variety of information and, in particular, information in which the information consumer may, or does, have an interest. Given the extent to which computers now permeate society, and particularly in view of the escalation of networking of those computers in various ways, there is increasing recognition of the capability of using computers, and, in particular, computers (and other devices) that are interconnected in a network, as an information dissemination tool that can satisfy the interests of both information providers and information consumers.

For example, information providers have used public computer networks (e.g., the Internet) and private computer networks (e.g., commercial online services such as America Online, Prodigy and CompuServe) to disseminate their information. This information can be displayed to a computer user having access to the network directly in response to a request from the user or indirectly (i.e., without request by the user) as a result of another action taken by the user. While these methods of information dissemination and acquisition can be effective, they do not exhaust the possibilities.

In a different vein, historically, computers have frequently included screen saving mechanisms ("screen savers") intended to prevent the phosphors of a computer display screen from burning out when the same image remains on the screen for a long period of time, such as might occur during a long period of inactivity while the computer is operating. As computer display screen technology has progressed, the use of screen savers to preserve the display screen has become increasingly unnecessary. However, the use of screen savers has continued—even proliferated—likely due to the aesthetic or entertainment value provided by the imagery of many screen savers. Further, the use of "wallpaper" (i.e., a pattern generated in the background portions on a computer display screen) in computer display screens has also arisen, largely one would suspect because of the aesthetic or entertainment value of the wallpaper imagery. While the use of screen savers and wallpaper with computer displays appeals to many users because of the imagery they present to the user, screen savers and wallpaper have not heretofore been used as a means to convey information from information providers to computer users. Further, screen savers and wallpaper have previously been implemented as relatively simple, self-contained computer application programs that are not typically integrated with other application programs or other aspects of computer operation. In particular, screen saver and wallpaper application programs have not been constructed to enable retrieval of display content from a remote location via a computer network.

2

SUMMARY OF THE INVENTION

An attention manager according to the invention presents information to a person in the vicinity of a display device in a manner that engages the peripheral attention of the person. Often, the display device is part of a broader apparatus (e.g., the display device of a computer). Generally, the attention manager makes use of "unused capacity" of the display device. For example, the information can be presented to the person while the apparatus (e.g., computer) is operating, but during inactive periods (i.e., when a user is not engaged in an intensive interaction with the apparatus). Or, the information can be presented to the person during active periods (i.e., when a user is engaged in an intensive interaction with the apparatus), but in an unobtrusive manner that does not distract the user from the primary interaction with the apparatus (e.g., the information is presented in areas of a display screen that are not used by displayed information associated with the primary interaction with the apparatus).

The information is embodied as one or more sets of content data. The sets of content data represent sensory data; typically, the sensory data is either video or audio data. Each set of content data is formulated by a content provider and made available for use by an attention manager according to the invention. Each content providing system can provide more than one set of content data. The content providing systems provide user interface tools that enable a particular set of content data to be requested. Once one or more sets of content data has been acquired, a content display system integrates scheduling information for all sets of content data to produce a schedule according to which an image or images corresponding to the sets of content data are displayed on a display device associated with the content display system.

A set or sets of instructions for enabling a display device to selectively display an image or images generated from a set of content data are also made available for use by the content display systems. Typically, the instructions enable images generated from content data to be displayed automatically, without user intervention, in a predetermined manner, thereby enhancing the capability of the invention to occupy the user's peripheral attention. Further, the attention manager can be implemented so that the instructions are automatically acquired (or updated, if necessary) each time a user requests acquisition of a set of content data, thereby making acquisition of the instructions transparent to the user of the attention manager and thus increasing the ease of use for the user. The instructions can include application instructions, control instructions and content data acquisition instructions. The application instructions can include operating instructions for beginning, managing, and terminating operation of the attention manager on a content display system, content display system scheduling instructions for scheduling the display of content data on a content display system, and installation instructions for installing the operating instructions and content display system scheduling instructions on a content display system. The control instructions can include display instructions for enabling generation of images from the content data on a particular type of display device or from a particular type of content data, and content data scheduling instructions for enabling temporal control of the display of the images generated from a set or sets of content data. The content data acquisition instructions can include acquisition instructions for enabling the acquisition of a set of content data, content data update instructions for enabling update of a previously acquired set of content data, and user interface installation instructions

6,034,652

3

for enabling provision of a user interface that allows a person to request a set of content data from a content providing system. Each of the application, control and content data acquisition instructions could be acquired from a content provider, or any one or all of the sets of instructions could be acquired from an application manager that provides generic sets of instructions that can be tailored as necessary or desirable by a content provider. Additionally, audit instructions can be made available that enable monitoring of usage of the attention manager.

According to one aspect of the invention, an attention manager engages the peripheral attention of a person in the vicinity of a display device of an apparatus by acquiring one or more sets of content data from a content providing system and selectively displaying on the display device, in an unobtrusive manner that does not distract a user of the apparatus from a primary interaction with the apparatus, an image or images generated from the set of content data. According to a further aspect of the invention, the selective display of the image or images begins automatically after detection of an idle period of predetermined duration (the "screen saver embodiment"). This aspect can be implemented, for example, using the screen saver API (application program interface) that is part of many operating systems. According to another further aspect of the invention, the selective display of an image or images occurs while the user is engaged in a primary interaction with the apparatus, which primary interaction can result in the display of an image or images in addition to the image or images generated from the set of content data (the "wallpaper embodiment"). If multitasking is allowed by the apparatus (e.g., by the computer operating system) with which the attention manager is used, the attention manager can be implemented so that, when operation of the attention manager is terminated, the user is returned to the state of the primary interaction that existed when operation of the attention manager began. The attention manager can also be implemented so that, during operation of the attention manager, the user is presented with a number of options regarding further use of the attention manager. In particular, one of the options can allow additional information to be obtained that is related to the set of content data for which an image is being displayed. Where the attention manager is implemented as part of a network, this option can enable information to be obtained from a remote information source via the network. Another option that can be implemented allows a user to specify a satisfaction level for a set of content data from which an image or images is being displayed, thereby affecting the frequency with which that set of content data is used by the attention manager in the future.

According to another aspect of the invention, an attention manager that engages the peripheral attention of a person in the vicinity of a display device includes a content display system associated with the display device, a mechanism that can communicate with the content display system via a first communications mechanism to provide to the content display system a set of instructions for enabling the display device to selectively display content data, and a content providing system that can communicate with the content display device via a second communications mechanism to provide a set of content data to the content display system. The content display system uses the provided set of instructions to selectively display on the display device an image or images generated from the provided content data. The attention manager according to this aspect of the invention can further include an application management system that

4

can communicate via a third communications mechanism to provide to either the content providing system or the content display system one or more sets of instructions for enabling a display device to selectively display an image or images generated from a set of content data. In the former case, the content providing system can, in turn, communicate with the content display system to provide the one or more sets of instructions. The attention manager according to this aspect of the invention can be implemented, for example, using existing computer networks of information sources, such as the Internet (in particular, the World Wide Web) or commercial online services, advantageously making use of pre-existing hardware and software for enabling communication over those networks. Typically, though not necessarily, an attention manager according to this aspect of the invention will include multiple content display systems and multiple content providing systems. The content providing systems will each be capable of providing one or more sets of content data, so that, overall, there will be multiple available sets of content data which can be of different types. There can also be multiple sets of instructions for enabling a display device to selectively display an image or images generated from a set of content data, which sets of instructions may be tailored to display images from particular types of content data or to display content data using a particular display device.

According to yet another aspect of the invention, a computer readable medium can be encoded with one or more computer programs for enabling acquisition of a set of content data and display of an image or images generated from the set of content data on a display device during operation of an attention manager. The instructions of the computer program can include: i) acquisition instructions for enabling acquisition of a set of content data from a specified information source, ii) user interface installation instructions for enabling provision of a user interface that allows a person to request the set of content data from the specified information source, iii) content data scheduling instructions for providing temporal constraints on the display of the image or images generated from the set of content data, and iv) display instructions for enabling display of the image or images generated from the set of content data. The computer readable medium can also further include content data update instructions for enabling acquisition of an updated set of content data from an information source that corresponds to a previously acquired set of content data, the content data update instructions specifying where and when to obtain the updated set of content data. The content data scheduling instructions can specify, for example, the duration of time that the image or images generated from a set of content data can be displayed, an order in which the images generated from a plurality of sets of content data are displayed, a time or times at which the image or images generated from a set of content data can or cannot be displayed, and/or constraint on the number of times that the image or images generated from a set of content data can be displayed. The display instructions can be tailored to enable display of the image or images generated from a set of content data on a display device of a particular type, or display of an image or images generated from a set of content data of a particular type.

According to still another aspect of the invention, a computer readable medium can be encoded with one or more computer programs for enabling a content display system to selectively display on a display device, in an unobtrusive manner that does not distract a person from a primary interaction with an apparatus associated with the display device, an image generated from a set of content data. The instructions of the computer program can include:

6,034,652

5

i) operating instructions for beginning, managing and terminating the selective display of the image on the display device, ii) content display system scheduling instructions for scheduling the display of the image on the display device, and iii) installation instructions for installing the operating instructions and content display system scheduling instructions on a content display system. The computer readable medium can also further include audit instructions for monitoring usage of the content display system to selectively display an image generated from a set of content data.

The attention manager according to the invention is a new and useful mechanism for providing information to users of the attention manager. The attention manager provides information in which a user has expressed an interest and, importantly, information that the user might not otherwise expend adequate energy to obtain. The user can tailor the information provided by interacting with specific information sources to indicate interest in particular information provided by a specific information source while the user is perusing other information from that information source (as opposed to giving a general indication of interest in information on a particular subject or of a particular kind, from which indication information that matches the indicated interest is automatically provided from various information sources). The user can also choose information from a wide variety of information sources; in particular, when the attention manager is implemented using a network (e.g., the Internet), the user can acquire information from a wide variety of remote information sources. Additionally, the information is presented to the user in a manner that uses portions of the user's attention capacity that may otherwise be unused or filled with extraneous information.

The attention manager according to the invention also provides a new and useful information dissemination tool to content providers. The attention manager affords an opportunity to content providers to disseminate their information to users that are interested in receiving such information, enabling the content providers to provide better directed information dissemination. Moreover, the attention manager provides access to the previously unused attention capacity of those interested users. Additionally, the attention manager allows content providers to tailor particular aspects of the attention manager as desired by the content provider, such as the acquisition of updated sets of the content provider's content data (e.g., the frequency of such updates), the display scheduling and manner of display of the content provider's content data, and the user interface that enables users to specify acquisition of the content provider's content data.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flow chart of a method that implements an attention manager according to an embodiment of the invention.

FIG. 2 is a block diagram of a system for implementing an attention manager according to an embodiment of the invention.

FIGS. 3A, 3B and 3C are schematic diagrams illustrating the functional components of an application manager, a content providing system and a content display system, respectively, according to an embodiment of the invention.

FIG. 4 is a flow chart of a method, according to an embodiment of the invention, for acquiring and updating sets of content data.

FIGS. 5 to 9C together are a flow chart of a method that implements an attention manager according to another embodiment of the invention.

6

FIG. 6 illustrates a computer display screen including a user interface, according to one embodiment of the invention, that can be used to enable a user to specify a control option.

DETAILED DESCRIPTION OF THE INVENTION

According to the invention, an attention manager presents information to a person in the vicinity of a display device (or devices) in a manner that engages at least the peripheral attention of the person. "Display device", as used herein, encompasses any device that presents sensory stimulus to the person and includes, for example, computer video display devices, televisions and audio speakers. Further, here, "in the vicinity of" means any location with respect to the display device from which the person can perceive the information being presented. For example, if the information is being presented in a visual form, then "in the vicinity of" means any location from which the person can see the information. Or, if the information is being presented in an aural form, then "in the vicinity of" means any location from which the person can hear the information.

Often, the display device is part of a broader apparatus that can be utilized by a user for a primary interaction that is unrelated to the attention manager. (However, the attention manager can also be used with a display device that is not part of a broader apparatus, the user engaging in a primary interaction with the display device.) For example, the display device can be part of a computer that can be used to implement any of a number of application programs (e.g., word processing programs, computer games, spreadsheets, etc.). The person whose attention is engaged by the attention manager can be the user or another person in the vicinity of the display device. In one embodiment of the invention, the information is presented by the attention manager while a primary interaction is ongoing, but during inactive periods (i.e., when the user is not engaged in an intensive interaction with the apparatus). In another embodiment of the invention, the information is presented by the attention manager during active periods (i.e., when the user is engaged in an intensive interaction with the apparatus), but in an unobtrusive manner that does not distract the user from the primary interaction (e.g., the information is presented in areas of a display screen that are not used by displayed information associated with the primary interaction). Generally, then, an attention manager according to the invention makes use of "unused capacity" of a display device, "unused capacity" being defined broadly to include, for example, the embodiments mentioned above, i.e., both temporal (e.g., the first-described embodiment above) and spatial (e.g., the second-described embodiment above) dimensions.

The information is embodied by one or more sets of content data. Each set of content data is formulated by a content provider and made available by a corresponding content providing system for use with the attention manager. Each content providing system can provide more than one set of content data. Moreover, each set of content data can include one or more "clips", each clip being a definable portion of the set of content data that is used to generate a particular "image." The term "image" is used broadly here to mean any sensory stimulus that is produced from the set of content data, including, for example, visual imagery (e.g., moving or still pictures, text, or numerical information) and audio imagery (i.e., sounds). The content providing systems can also provide user interface tools that allow a user of the attention manager to specify that they want to obtain a particular set of content data. Once obtained, one or more

6,034,632

7

images generated from the clips of one or more sets of content data are displayed by a content display system. The content display system integrates scheduling information associated with the sets of content data to produce a schedule according to which the images corresponding to the sets of content data are displayed for a particular user of the attention manager.

A set or sets of instructions for enabling a display device to selectively display images generated from one or more sets of content data are also made available to users of the attention manager. The instructions include application instructions, control instructions and content data acquisition instructions. Typically, the instructions enable images generated from content data to be displayed automatically, without user intervention, in a predetermined manner, thereby enhancing the capability of the attention manager to occupy the user's peripheral attention. Different sets of instructions can be formulated, such that only images generated from sets of content data that are compatible with a particular set of instructions can be displayed using that set of instructions. Typically, an application manager establishes a standard set or sets of instructions which content providers can tailor to fit their needs or desires.

As indicated above, the sets of content data represent sensory data, i.e., data that can be used to generate images as defined above. Typically, the sensory data is either video or audio data. The kinds of content data that can be used with the attention manager are virtually limitless. For example, video data that might be used as content data includes data that can be used to generate advertisements of interest to the user, moving and still video images which can be real-time or pre-recorded (e.g., nature scenes, pictures of family members, MTV music segments, or video from a camera monitoring a specified location, such as ski slopes or a traffic intersection, for conditions at that location), financial data (e.g., stock ticker information) or news summaries. Audio data that might be used as content data includes data that can be used to generate, for example, music or news programs (e.g., radio talk shows).

The attention manager according to the invention is useful both to users of the attention manager and to content providers. For users, the attention manager provides information to a user in which the user has expressed an interest. In particular, the attention manager provides information to a user that the user might not otherwise expend adequate energy to obtain. Additionally, the information is presented to the user in a manner that uses portions of the user's attention capacity that may otherwise be filled with extraneous information. Further, a variety of information can be displayed (i.e., images can be generated from more than one set of content data), so that the user does not have to choose particular information to the exclusion of all other information.

For content providers, the attention manager affords an opportunity to disseminate information to users that are interested in receiving such information, thus enabling the content providers to provide better directed information dissemination. Moreover, the attention manager provides access to the previously unused attention capacity of those interested users. Further, since information from more than one content provider can be displayed, content providers are more likely to have their information displayed, since their information is displayed in addition to, rather than instead of, the information of other content providers, thereby reducing the need to compete with other content providers for the attention of the user.

FIG. 1 is a flow chart of a method 100 that implements an attention manager according to an embodiment of the inven-

8

tion. The method 100 is performed by a content display system according to the invention. The content display system can be implemented, for example, using a digital computer that includes a display device and that is programmed to perform the functions of the method 100, as described below. Below, the method 100 is described as implemented on such a digital computer, though the method 100 could be implemented on other apparatus.

As shown by block 101, initially (i.e., before operation of the attention manager begins), a user is engaged in a primary user interaction, e.g., a primary user interaction with a computer. Though shown in FIG. 1, the primary user interaction of block 101 does not form part of the method 100 according to the invention. "Primary user interaction" is to be construed broadly and, generally, includes any operation of the computer (or other apparatus with which the user is engaging in an interaction) other than operation that is part of the attention manager according to the invention. When the user is interacting with a computer, the primary user interaction includes any operation of the computer that occurs to enable or to support the performance of the function or functions that provide the basis for the user's use of the computer. For example, the primary user interaction can be the use of any of a variety of conventional application programs (e.g., word processing programs, spreadsheet programs, personal finance programs, game programs, drawing programs, online services and Web browsers, among others). The primary user interaction can also be, for example, simply the operation of a conventional computer operating system, such as the Windows (e.g., Windows 3.1, Windows NT or Windows 95) or DOS operating systems produced by Microsoft Corp. of Redmond, Wash. or the Macintosh operating system produced by Apple Computer, Inc. of Cupertino, Calif., among others. While, typically, the display device produces a display as a result of the primary user interaction, this need not necessarily be the case.

The method 100 actually begins with the block 102. In the step shown in the block 102 (referred to hereinafter as step 102), a determination is made as to whether an "idle period" has occurred. Generally, as used herein, "idle period" refers to a period of time of specified duration during which a specified condition does not occur. However, typically, the specified condition is one having the characteristic that failure of the condition to occur is indicative of an extended lack of intensive (or focused) interaction with the computer by the user ("user inactivity"). For example, the specified condition could be the lack of an input from an input device of the computer, e.g., the absence of striking a key on a keyboard, clicking a mouse, pressing on a touch-sensitive area of a touchscreen or issuing a voice command. Alternatively, the attention manager could be implemented with an apparatus that can monitor the environment of the apparatus (e.g., with a video camera) and evaluate the environment to ascertain that an "idle condition" (e.g., the viewing direction of the user of the apparatus is turned away from the apparatus by a specified amount for a specified period of time) has occurred, such idle condition triggering operation of the attention manager.

Theoretically, any duration of time can be specified to define the idle period. However, practically, the duration of time necessary to constitute an idle period cannot be so short that the attention manager begins operating at times that inhibit the user's primary interaction with the computer or that distract or annoy the user. Further, the duration of time chosen, as indicated above, should be sufficiently long to indicate an extended lack of interaction with the computer, suggesting that the user is not engaged in an interaction with

6,034,652

9

the computer that the user would not want to have interrupted. However, the duration of time should not be so long that, for periods of user inactivity of a typical duration, the amount of time that the attention manager operates is undesirably short. In sum, choosing the duration of time that defines an idle period involves a balancing of the above considerations. Illustratively, the idle period can be defined as a period of between thirty seconds and two minutes during which the specified condition (e.g., user interaction with an input device) does not occur.

While detection of the idle period can be implemented in any suitable manner, one way in which such detection can be implemented is by monitoring an idle timer that is part of a screen saver API (application program interface) that is, in turn, part of an operating system used to operate the computer. Such screen saver APIs are commonly found in current operating systems such as the Windows or Macintosh operating systems discussed above. The idle timer could be monitored and a signal that an idle period has occurred generated when the magnitude of the idle time as indicated by the idle timer reaches a predefined threshold.

Detection of an idle period as the basis for beginning operation of the attention manager is an indirect activation of the attention manager. In an alternative embodiment, step 102 of the method 100 is modified so that the attention manager is activated directly by the user. In other words, step 102 would consist of waiting for explicit direction from the user to begin operation of the attention manager. Such explicit direction could be enabled with an appropriate user interface, such as an on-screen icon or a menu selection, that is always present on the display screen of the display device as part of a standard interface that is provided by the operating system. Examples of such standard interfaces are the "Apple Menu" provided as part of the Macintosh operating system, and the "Start Menu" or desktop icons provided as part of the Windows 95 operating system.

Returning to FIG. 1, if, in step 102, an idle period has not occurred, then the primary user interaction continues (block 101). The method 100 continues executing the step 102 at predefined time intervals (typically very short time intervals), thereby continually and frequently checking for the occurrence of an idle period.

If, in step 102, an idle period is detected, then, in the step shown in the block 103 (hereinafter referred to as step 103), a determination is made as to whether there are any sets of content data available for use in generating a display. (Hereinafter, reference is sometimes made to "displaying content data" or "displaying a set of content data"; it is to be understood that this means displaying images generated using the content data or set of content data.) Herein, "content data" refers to data that is used by the attention manager to generate displays (e.g., video images or sounds, or related sequences of video images or sounds). A "set of content data" refers to a related set of such data that is used to generate a particular display. A "clip" refers to a definable portion of a set of content data that is used to generate a particular image; a set of content data can include one or more clips and, therefore, can be used to generate one or more images. The acquisition of content data by the content display system is described in more detail below. Here, it is sufficient to note that, over time, an attention manager can acquire any number of sets of content data that can be displayed by the content display system.

If, in step 103, no sets of content data are available for display, then the primary user interaction continues (block 101). The method 100 continues executing the steps 102 and

10

103 at predefined time intervals, continually checking for the occurrence of an idle period and the acquisition of at least one set of content data.

If, in step 103, at least one set of content data is available for display, then, in the step shown in the block 104 (hereinafter referred to as step 104), the available sets of content data are scheduled for display by the content display system. (Alternatively, in other embodiments of the invention, scheduling of the sets of content data can occur before the method 100 begins. Such scheduling might be implemented, for example, so that each time a new set of content data is received by the content display system, the schedule is revised to include the new set of content data.) Typically, when the content display system acquires a new (or updated) set of content data, scheduling information for that set of content data is also acquired. Taken together, the scheduling information for all of the sets of content data is used to determine a schedule for display of the sets of content data by the content display system. Generally, determining a display schedule involves specifying the order in which the sets of content data are to be displayed and the duration of time for which each set of content data is to be displayed. The determination of the display schedule can also accommodate (to the extent possible) any special scheduling parameters for particular sets of content data (e.g., restrictions specifying when a particular set of content data must be displayed or cannot be displayed), mediating any conflicts between the display requirements of particular sets of content data. Often, though not necessarily, once the order and duration of display are established, the sets of content data are repetitively displayed by cycling through the display schedule repeatedly until operation of the attention manager is terminated. However, even where such iteration through the display schedule occurs, the display schedule can also accommodate scheduling parameters that delete sets of content data from the display schedule during particular iterations, thereby, for example, controlling the frequency with which particular sets of content data are displayed. The display schedule can be stored in an appropriately structured database, as known by those skilled in the art, that is stored in a memory of the computer used to implement the content display system.

Any appropriate set of rules, that can, for example, be arranged in any appropriate hierarchical manner, can be used for establishing a display schedule and, in particular, mediating conflicts between conflicting scheduling parameters associated with different sets of content data. For example, one rule for mediating conflicts may give preference to displaying sets of content data so that the sets of content data are displayed inversely to the order in which they were obtained by the content display system. This rule might be further specified so that a set of content data that has never previously been displayed by the attention manager is displayed prior to display of a set of content data that has been previously displayed, even though an update of the previously displayed set of content data has been obtained at a later time than that at which the never displayed set of content data was obtained. Another rule for mediating conflicts might resolve a conflict between two sets of content data having scheduling parameters that specify display at the same sequential position in the display schedule by randomly selecting one of the sets of content data to be displayed first during each iteration through the display schedule. Still another rule for mediating conflicts might establish a hierarchy of kinds of content data, with sets of content data of kinds at the top of the hierarchy being given preference for display over those at the bottom. Yet another

6,034,652

11

rule or set of rules for mediating conflicts may involve performing some sort of analysis of the characteristics of the sets of content data that have been obtained by a particular content display system to ascertain preferences indicated thereby, and giving preference to sets of content data that are evaluated to be relatively more preferred. Scheduling rules of this kind would typically be part of the scheduling parameters provided independent of the content providers (i.e., in content display system scheduling instructions, as discussed elsewhere herein and, in particular, with respect to FIGS. 3A through 3C below).

Other scheduling rules, not directed to mediating conflicts between sets of content data, can also be used in determining a schedule. For example, any set of content data that has been initially obtained before a certain time and/or that has been last updated before a certain time (i.e., a set of content data that is "stale") can be automatically precluded from being inserted into the display schedule. This exclusion could further be restricted to apply only to certain sets of content data or content data of certain kinds. Similarly, the frequency with which a particular set of content data appears in a display schedule can be based upon how stale the set of content data is. Scheduling rules of this kind would typically be part of the scheduling parameters provided by a content provider for a set of content data (i.e., in tailored content data scheduling instructions, as discussed elsewhere herein and, in particular, with respect to FIGS. 3A through 3C below).

The particular scheduling rules used may be influenced by the characteristics of a particular embodiment of the attention manager, such as the available kinds of content data or the characteristics of the potential users of the attention manager. The particular scheduling rules used may also be influenced by the need or desire to simplify implementation of the scheduling rules.

Returning to FIG. 1, once the sets of content data have been scheduled for display, then, in the step shown in the block 100 (hereinafter referred to as step 100), a set of content data is displayed. The content display system is provided with one or more sets of display instructions to enable display of the set or sets of content data on the display device (as discussed elsewhere herein and, in particular, with respect to FIGS. 3A through 3C below).

After a set of content data has been displayed, then, in the step shown in the block 106 (hereinafter referred to as step 106), a determination is made as to whether operation of the attention manager has been terminated. Generally, operation of the attention manager can be terminated either directly or indirectly. Indirect termination of operation of the attention manager can be effected by, for example, causing operation of the attention manager to terminate when the specified condition (the non-occurrence of which is used to signal an idle period) occurs. For example, the attention manager can be terminated if the user makes an input to the computer using an input device, e.g., strikes a key on a keyboard, clicks a mouse, presses on a touch-sensitive area of a touchscreen or issues a voice command. For indirect termination, it may be desirable to add a further step or steps to the method 100 that, upon an indication that indirect termination should occur (e.g., the occurrence of the specified condition), asks the user to confirm that termination of the attention manager is, in fact, desired, and, if so, terminates the attention manager upon appropriate specified user input. In contrast to indirect termination, direct termination of operation of the attention manager can be effected by, for example, causing operation of the attention manager to terminate when the user selects a control option that specifies such termination, as described in more detail below with respect to FIGS. 3A, 3B and 6.

12

If, in step 106, operation of the attention manager has been terminated, then the primary user interaction begins again (block 101). The method 100 then begins executing the step 102 again, checking for the occurrence of an idle period.

If, in step 106, operation of the attention manager has not been terminated, then, in the step shown in the block 107 (hereinafter referred to as step 107), a determination is made as to whether there is an additional set of content data to be displayed. Typically, in operation of an attention manager according to the invention, there will always be another set of content data to be displayed, since, as discussed above, the sets of content data in the display schedule are iteratively displayed until operation of the attention manager is terminated. However, this need not be the case. For example, a limit can be established on the number of times that each set of content data can be displayed, or on the total number of times that any set of content data is displayed.

If, in step 107, there are no additional sets of content data to be displayed, then the primary user interaction begins again (block 101). The method 100 then begins executing the step 102 again, checking for the occurrence of an idle period.

If, in step 107, there are additional sets of content data to be displayed, then the method 100 returns to the step 102 and displays a set of content data in accordance with the previously determined display schedule. Steps 102, 106 and 107 are continuously performed, resulting in the continuous display of sets of content data, until either the user terminates the attention manager (step 106) or there are no more sets of content data to be displayed (step 107).

In another embodiment of the invention, a step could be added to the method 100, either in place of or in addition to the step 107, or as part of the step 106, that causes operation of the attention manager to terminate after the attention manager has been operating for a specified period of time.

Further, in another embodiment of the invention, an appropriate step or steps could be added to the method 100 so that, at a specified time, such as after each iteration through the display schedule, the method 100 returns to the step 104 and re-determines the display schedule.

As described above, when the method 100 ends, the primary user interaction (block 101) begins again. Preferably, the primary user interaction begins again with the status existent at the time that the method 100 began. Thus, the primary user interaction must be held in abeyance while the method 100 is operating. This can be accomplished by implementing the method 100 (or any other embodiment of the attention manager) with a content display system that is implemented on a computer that operates with an operating system that allows "multi-tasking" (here used to mean either the suspension of one program while one or more other programs operate, or the execution of one program simultaneously with the execution of one or more other programs). The Windows and Macintosh operating systems (mentioned above), among others, are operating systems having this characteristic. Where the attention manager is implemented using a screen saver API that is part of the operating system, such multitasking occurs automatically as a characteristic of the screen saver API, i.e., when operation of the attention manager ends, the user is returned to the status of the primary interaction existent at the time that the attention manager started operating. In multitasking operating systems that do not include a screen saver API, this feature of the invention can be implemented by use of an appropriately programmed device driver, as known by

6,034,652

13

those skilled in the art, that monitors user interaction, suspending and restarting the primary user interaction at the beginning and end of operation of the attention manager.

The method 100 (FIG. 1) described above is an embodiment of the invention in which the attention manager presents information to a person (which can be the user or another person) in the vicinity of the display device during inactive periods when a user is not engaged in an intensive interaction with the computer (as indicated by the step 102 which checks for the occurrence of an "idle period" before beginning operation of the attention manager). As indicated above, in other embodiments of the invention, the attention manager presents information to the person during active periods, but in an unobtrusive manner. In such embodiments, video content data could be presented, for example, as "wallpaper" on the display screen of a video display monitor. Audio content data according to these embodiments could be presented in the same way as for the embodiments of the method 100 described above. For implementation of such embodiments of the invention, the step 102 of the method 100 could be modified to be a determination as to whether the attention manager has been activated (typically this would require direct activation by the user). Alternatively, step 102 could be eliminated altogether and the attention manager could be implemented to operate at any time that the computer is operating and sets of content data are available for display (step 103). For these embodiments, it is, as above, necessary that the content display system be implemented on a computer operated by an operating system that allows multi-tasking as described above. In particular, simultaneous operation of programs must be allowed, since the attention manager operates while the primary user interaction is ongoing (note that the relationships between the block 101 and the method 100 shown in FIG. 1 are not present in these embodiments of the invention).

Though not confined to such use, the attention manager according to the invention is envisioned as having particular use as a system implemented on, and used by, a network of computers. In such an implementation, each content providing system is implemented on a content provider computer. (It is possible to have more than one content providing system on a content provider computer.) Content display systems are implemented on user computers. The content provider computers and user computers are integrated together into a network such that each user computer can communicate with one or more of the content provider computers. The content provider computers need not (and typically would not) communicate with each other. Likewise, the user computers need not (and typically would not) communicate with each other. Further, each user computer need not communicate with all, or even more than one, of the content provider computers. For example, an attention manager according to the invention could be implemented so as to make use of a network such as the Internet. In particular, the graphical attributes of the World Wide Web would be particularly useful in enabling the provision of user interfaces that allow users to access the attention manager while visiting network sites of content providing systems.

FIG. 2 is a block diagram of a system 200 for implementing an attention manager according to an embodiment of the invention. The system 200 includes an application manager 201, a multiplicity of content providing systems, shown as Content Providers 1 through n (content providing systems 202a, 202b and 202c are illustrated in FIG. 2), and a multiplicity of content display systems, shown as Users 1

14

through n (content display systems 203a, 203b and 203c are illustrated in FIG. 2). Hereinafter, the content providing systems and content display systems are referred to generally using the numerals 202 and 203, respectively. In FIG. 2, the solid lines indicate that communication must occur in the system 200 and the dashed lines indicate that communication may occur. However, note that, in another embodiment of the invention, the application manager 201 is not present, and communication between any particular content display system and particular content providing system need not necessarily occur.

The application manager 201, content providing systems 202 and content display systems 203 can be implemented using appropriately programmed digital computers. Generally, the computers can be any conventional digital computers including an input device (such as a keyboard, mouse or touch screen), an output device (such as a conventional computer display monitor and/or one or more audio speakers), a processing device (such as a conventional microprocessor), a memory (such as a hard disk and/or random access memory), additional conventional devices necessary to interconnect and enable communication between the above-listed devices, and communications devices (e.g., a modem) for enabling communication with other computers of the system. For example, the application manager 201 and content providing systems 202 can be implemented using conventional server computers, while the content display systems 203 can be implemented using conventional client computers. The application manager 201, content providing system 202 and content display systems 203 could also themselves each be implemented by a client-server network of computers. Communication between the computers can be accomplished using any appropriate communication transmission lines, such as conventional telephone lines, or high speed data transmission systems such as T1, T3 or ISDN. The communication can be managed using any appropriate conventional networking methods (e.g., computer programs and protocols) and apparatus, as known by those skilled in the art. In particular, as described further below, the computers are programmed to enable the content display systems 203 to communicate with the content providing systems 202 and application manager 201 even without direct action by the user. In addition to being programmed to enable networking, each of the computers is also appropriately programmed, as described above and below, to perform the functions of the application manager 201, content providing systems 202 and content display systems 203, as appropriate.

FIGS. 3A, 3B and 3C are schematic diagrams illustrating the functional components of the application manager 201, a content providing system 202 and a content display system 203, respectively, according to an embodiment of the invention. Each of the functional components are represented by a set of instructions and/or data. (In particular, each of the sets of instructions may include, if appropriate, data related to accomplishment of the functions associated with the set of instructions; similarly, a set of content data may include, if appropriate, instructions that enable generation of an image from the set of content data.) Each of these sets of instructions and/or data can be embodied in an appropriate computer program or set of computer instructions (the latter capable of including computer instructions and data), or an appropriate set of data configured for use by a set or sets of instructions (e.g., computer program) that must interact with the set of data in order to implement the attention manager.

The application manager 201 stores a variety of instructions for use in implementing the attention manager. As

6,034,652

15

shown in FIG. 3A, generally, the application manager 201 stores application instructions 310, control instructions 320, and content data acquisition instructions 330 that can be disseminated to the content display systems 203 and content providing systems 202 as necessary or appropriate. The application manager 201 can also store audit instructions 340 that can be used to enable monitoring of usage of the attention manager.

The application instructions 310 include operating instructions 311 for beginning, managing and terminating operation of the attention manager on a content display system 203, and content display system scheduling instructions 313 for scheduling the display of content data on a content display system 203. The method 100 (FIG. 1) described above is one embodiment of such application instructions 310. The application instructions 310 also include installation instructions 313 that enable the other instructions used by the attention manager to be implemented using the hardware that is part of and associated with a particular content display system 203. The installation instructions 313 can be implemented as known by those skilled in the art. For example, the installation instructions 313 can be a "plug-in" or "helper" application program (such as a helper application that can be used with the Navigator and Mosaic software programs made by Netscape Communications Corp. of Mountain View, Calif.) that is used to process instructions or data of a particular type—in this case, instructions to implement the attention manager, and content data for use with the attention manager, that can be installed on the hardware of a particular content display system 203. There can be a multiplicity of such helper applications, each capable of operating on particular hardware that could be used to implement a content display system 203. The helper application enables the software program (i.e., Navigator or Mosaic) being used to access the sites of content providers to process references (e.g., Universal Resource Locators, or "URLs") to the particular type of instructions and/or data, so that sets of content data (including updated sets of content data) and the application instructions 310, control instructions 320 and content data acquisition instructions 330 (including updated versions of those instructions) can be acquired.

The control instructions 320 include display instructions 321 and content data scheduling instructions 322, as described in more detail below, that are typically enhanced by content providers in a particular manner that is appropriate for the content data that the content providers provide. The application manager 201 can (and typically does) store and disseminate multiple distinct sets of control instructions 320. Generally, the display instructions 321 of a particular set of control instructions 320 enable display of content data on a particular type of display device (e.g., a particular type of computer video display or a particular type of audio speaker) or display of a particular type of content data. Display instructions 321 that can be used with a particular display device are typically already developed by third parties (e.g., the maker of the display device) and are readily available. Tailoring of the display instructions 321 to display particular types of content data (such as instructions for displaying content data that is in the GIF format or the format of AutoDesk Animator FLC files) can be done by either the application manager 201 or a content provider. The content data scheduling instructions 322 provide temporal constraints on the display of particular sets of content data. As stored by the application manager 201, the content data scheduling instructions 322 are usually the same for each set of control instructions 320 and provide a generic set of scheduling instructions that can be tailored by a content provider.

16

The content data acquisition instructions 330 include acquisition instructions 331 for enabling the initial acquisition of a set of content data and instructions for implementing the attention manager, and content data update instructions 333 for enabling update of previously obtained sets of content data and attention manager instructions. The acquisition instructions 331 and content data update instructions 333 are generic sets of instructions that can be tailored by a content provider. The content data acquisition instructions 330 can also include user interface installation instructions 333 that enable content providers to install a user interface in the information environment (e.g., Web page) of the content provider so that users can request sets of content data from the content provider. Such user interface installation instructions are conventional and readily available for use with the attention manager of the invention.

As shown in FIG. 3B, the content providing systems 202 store one or more sets of content data 300 that can be disseminated to content display systems 203 as requested. The content providing systems 202 can also store the application instructions 310, control instructions 320, and content data acquisition instructions 330 described above.

As indicated above, each set of content data 300 defines a related group of data that is used to generate a particular display and includes one or more clips that each represent a definable portion of the set of content data that is used to generate a particular image. The content data 300 represents sensory data and can be, for example, video or audio data. A particular set of content data 300 can be formulated in different versions that are each compatible with content display systems 203 having particular characteristics. In particular, the characteristics of the display device of a content display system 203 can affect the formulation of a set of content data 300. For example, for computer video display monitors, the formulation of a version of a set of content data 300 can depend on the size of the display screen (e.g., horizontal length by vertical length), the display resolution (e.g., the number of horizontal pixels by the number of vertical pixels), the color depth (number, e.g., 256, of possible colors) and the characteristics of the display drivers for the display device. The formulation of a version of a set of content data 300 could also depend upon the operating system being used by the computer on which the content display system 203 is implemented or other characteristics of the computer, such as the speed with which the display device can be operated (insofar as that speed is affected by the characteristics of the computer such as processor speed). Generally, a set of content data 300 can be formatted as known by those skilled in the art in view of the above considerations.

As indicated above, the control instructions 320 (as well as the content data acquisition instructions 330) are typically enhanced by content providers as appropriate for particular content data. The manner in which these instructions can be tailored by content providers is desirably required to conform to a specified format. Below, a description is given of package files that can be used for tailoring the control instructions 320 and content data acquisition instructions 330, as well identifying the location of content data. These package files can be created using an appropriate computer program (package file editor) that can be provided by, for example, the application manager 201 and that enables this tailoring to be accomplished easily and according to the specified format.

The content provider can tailor the content data scheduling instructions 322 to indicate the duration of time that a particular set of content data can be displayed ("duration

6,034,652

17

instructions"). Generally, the duration instructions can be arbitrarily complex and can vary in accordance with a variety of factors, including, for example, the particular time at which the set of content data 300 is displayed after the attention manager begins operating, or the number of previous times that the set of content data 300 has been displayed during a continuous operation of the attention manager. The content provider can also tailor the content data scheduling instructions 323 to indicate an order in which the clips of a set of content data 300 are displayed, as well as the duration of the display for each clip ("sequencing instructions"). The content provider can also tailor the content data scheduling instructions 323 to indicate particular times or ranges of times at which a set of content data 300 can or cannot be displayed ("timing instructions"). These times can be absolute (e.g., a particular clock time on a particular day, a particular day or days during a week, after or before a specified date) or relative (e.g., not before or after a specified duration of time since the attention manager began operation, first or not first among the sets of content data 300 to be displayed, not after a particular kind or set of content data 300). The content provider can also tailor the content data scheduling instructions 323 to specify a maximum number of times that the set of content data 300 can be displayed after the attention manager begins operating or a maximum number of times that the set of content data 300 can be displayed over any number of operations of the attention manager ("saturation instructions").

The content provider can also tailor the display instructions 321 to display a particular set or sets of content data. The display instructions 321 can be tailored, for example, according to the type or types of the content data. The type of content data indicates the manner in which an image or images are generated from the content data (i.e., how the bit patterns in a particular clip are transformed into an image). The type of content data is typically established as a consequence of the manner (e.g., with a particular software application program such as the Photoshop or Premiere programs produced by Adobe Systems of Mountain View, Calif.) in which a particular clip is created. The installation instructions 313, discussed above, enable content data of different types to be obtained by the attention manager. Generally, the possible types of content data can be confined to an enumerated set of standard data types, such as the Mime data types used with the World Wide Web. As will be more readily understood from the description below, the type of content data can be specified, for example, in a field of the clip part of a package file.

The ability to tailor sets of content data 300 and associated control instructions 320 for particular content display systems 203, before the sets of content data 300 and control instructions 320 are provided to those content display systems 203, is advantageous because it allows the tailoring to be done once, by the content providing system 202 or the application manager 301, rather than multiple times, once by each content display system 203 that uses the set of content data 300 and associated control instructions 320.

The content data acquisition instructions 330—in particular, the content data update instructions 333—are also tailored by content providers as appropriate for particular sets of content data 300. In particular, the content provider can tailor the content data acquisition instructions 330 to indicate where and when to obtain an updated set of content data 300. For example, the indication of where to obtain an updated set of content data 300 can be accomplished by specifying an appropriate network address of a content providing system 203. The network addresses can be speci-

18

fied by, for example, a URL used to identify, for example, an HTML file, an applet (a short application program written in Java or other suitable programming language), a script based on CGI or other suitable mechanism, or any other resource (i.e., computer program or set of data). The indication of when to obtain an updated set of content data 300 can be accomplished by specifying a time or times, either absolute time or times (i.e., particular dates and times during the day) or relative time or times (e.g., one month after the last acquisition/update of the set of content data 300). For example, the update schedule could be established to obtain updates every hour, every day or every week. Or, the update schedule could be established to obtain updates upon the occurrence of a particular event, such as a specified percentage increase or decrease in a stock market index. In general, the particular update schedule used will depend upon the character of the content data with which the update schedule is associated, e.g., content data representing stock prices would probably be updated more frequently than content data representing an advertisement.

As shown in FIG. 3C, the content display systems 203 store the application instructions 310, control instructions 320, and content data acquisition instructions 330 described above. The application instructions 310 use the control instructions 320 to display sets of content data 300 that are obtained (and updated, if appropriate) by the content data acquisition instructions 330. The application instructions 310 and control instructions 320 are discussed generally, and with respect to particular embodiments, in more detail above, while an embodiment of the content data acquisition instructions 330 is described below.

FIG. 4 is a flow chart of a method 400 according to the invention for acquiring and updating sets of content data, i.e., the method 400 is an embodiment, at least in part, of the acquisition instructions 331 and content data update instructions 333 of the content data acquisition instructions 330 discussed above with respect to FIGS. 3A through 3C. In the method 400, the steps shown by blocks 403 through 407 can be implemented in the acquisition instructions 331 and the steps shown by blocks 409 through 410 can be implemented in the content data update instructions 333. Generally, the steps of the method 400 can be implemented on an appropriately programmed digital computer that is programmed to perform the functions of the method 400, as described below. Below, the method 400 is described as implemented on such a digital computer, though the method 400 is not limited to such an implementation. The method 400 necessitates communication between a content display system 203 and one or more content providing systems 202. As will be understood by those skilled in the art of digital computer programming for computer network communications, when the method 400 is implemented using a programmed digital computer, particular steps of the method 400 could be implemented on either a content display system 203 or a content providing system 202.

In the step shown in the block 401 (referred to hereinafter as step 401), a set of content data is selected for display by the attention manager. Initially, in step 401, particular sets of content data are obtained as a result of direct request by the user. Any appropriate user interface can be used for enabling a user to directly request a particular set of content data. For example, Web pages on the World Wide Web could include graphical buttons for enabling users that visit the Web page to request particular sets of content data. Selection of a button on a Web page results in an indication to the appropriate content providing system 203 that the requesting content display system 203 has requested the set of content

6,034,652

19

data corresponding to the selected button to be transferred to the content display system 203. The user interface instructions 333 discussed above, that can be provided to each content providing system 202, can be used to create the user interface.

Selection of a set of content data in step 401 causes a set of acquisition instructions 331 to be transferred to the content display system 203. The acquisition instructions 331 include information identifying the site from which the set of content data can be obtained, as well as the site or sites from which instructions (e.g., application instructions 310, control instructions 330, content data acquisition instructions 330 and audit instructions 340) for implementing the attention manager can be obtained. These sites can be the same or different sites. The sites can be identified by, for example, using URLs, as described above. The acquisition instructions 331 can also include instructions for establishing an appropriate user interface (e.g., a desktop icon) in the content display system 203 that enables a user to cause the installation instructions 313 to be executed, thereby installing the attention manager in the content display system 203.

In the step shown in the block 403 (referred to hereinafter as step 403), a determination is made as to whether the requesting content display system 203 has the application instructions 310 (FIGS. 3A through 3C) that enable operation of the attention manager and scheduling of sets of content data 300. If the content display system 203 does not have the application instructions 310, then, in the step shown in the block 405 (referred to hereinafter as step 405), the content display system 203 uses the appropriate site identification provided by the content providing system 202 to obtain a version of the application instructions 310 (typically the most current version of the application instructions 310 that is compatible with the set of content data 300 requested by the user). The application instructions 310 can be provided by the content providing system 202 from which the set of content data 300 is being obtained. Alternatively, the application instructions 310 can be provided directly to the content display system 203 by the application manager 201 (or from some site other than a content provider or the application manager 201) by causing an appropriate instruction to be issued to the application manager 201 (or other site) by either the content providing system 202 or the content display system 203.

If the content display system 203 does have the application instructions 310 (step 403), then, in the step shown in the block 403 (referred to hereinafter as step 403), a determination is made as to which version or versions of the application instructions 310 the content display system 203 has. As indicated elsewhere, a particular set of content data 300 can (and typically will) be updated from time to time, thereby creating different versions of the set of content data 300. Likewise, the application instructions 310 can also be updated, thereby creating different versions of the application instructions 310. In general, a set of content data 300 can be updated without regard to whether the set of content data is compatible with all versions of the application instructions 310 (though the set of content data 300 must be compatible with at least one version of the application instructions 310). Likewise, the application instructions 310 can be updated without regard to whether any particular set of content data 300 is compatible with that version of the application instructions 310. Moreover, particular versions of the application instructions 310 may be compatible only with sets of content data 300 of certain types. Consequently, a particular content display system 203, even though the content display system 203 has the application instructions

20

310, may not have a version of the application instructions 310 that is compatible with the type and/or version of the set of content data 300 being requested.

It is necessary, therefore, to determine whether the content display system 203 has a version of the application instructions 310 that is compatible with the type and version of the set of content data 300 being requested so that, if necessary, a compatible set of application instructions 310 can be provided to the content display system 203. In the step shown in the block 404 (referred to hereinafter as step 404), this determination is made. If the content display system 203 does not have a compatible version of the application instructions 310, then, in step 405, the content providing system 202 (or, for example, the application manager 201) provides to the content display system 203 a version of the application instructions 310 (typically the most current version) that is compatible with the requested set of content data 300.

Alternatively, in step 404, a determination could be made as to whether the version of the application instructions 310 that the content display system 203 has is the most current version of a set of compatible application instructions 310. If the version is not the current version, then the content providing system 202 provides the current version (step 405), even if the version that the content display system 203 already had is compatible with the newly acquired set of content data 300.

Preferably, updated sets of application instructions 310 are made downwardly compatible with previous sets of application instructions 310, so that the updated application instructions 310 can be used with any previously obtained sets of content data that are compatible with a previous set of application instructions 310. If downward compatibility is not maintained, the updated set of application instructions 310 can replace a previous set of application instructions 310 and incompatible sets of content data can be removed from the schedule of sets of content data to be displayed (this can be accomplished by the use of appropriate instructions in the content display system scheduling instructions 312 that check for the compatibility of sets of content data with the existing set of application instructions 310) when the attention manager is operating. The content data update instructions 332 can also include instructions that ascertain the current version of the application instructions 310 and, for each set of content data 300 that is incompatible with the current version of the application instructions 310, seek to obtain, at the time scheduled for an update, an updated set of content data 300 that is compatible with the current version of the application instructions 310.

The steps 403 through 405 are advantageous in that they result in the provision of application instructions 310 to a content display system 203 only when such instructions are needed, thus minimizing the number of sets of application instructions that are made available. The steps 403 through 405 also minimize the amount of information that must be transmitted over communication lines to the content display system 203, thereby freeing those lines for other communication and minimizing the cost (i.e., cost of using the communication lines) associated with using the attention manager of the invention.

Returning to FIG. 4, in the step shown in the block 406 (referred to hereinafter as step 406), the content providing system 202 provides the current set of content data 300 to the content display system 203. (In practice, the set of content data 300 can be provided before, after or simultaneously with provision of the application instructions 310.)

6,034,652

21

Further, as described above, a particular set of content data 300 can exist in different versions that are each compatible with the content display system 203 to which the version of the set of content data 300 is being provided. The step 406 can include a determination as to the version or versions of the set of content data 300 that can be used by the requesting content display system 203, so that a properly formulated set of content data 300 is acquired.

A set of control instructions 320 and content data acquisition instructions 330 (FIGS. 3A through 3C) associated with the set of content data 300 can also be provided, as shown by the step of block 407 (referred to hereinafter as step 407). Typically, a check is made (like that for the application instructions 310 and providing similar benefits) to determine whether the content display system 203 already has a compatible (and/or current) version of the control instructions 320 and/or the content data acquisition instructions 330 associated with the set of content data 300 being obtained.

Each set of content obtained by a content display system 203 can be stored in a database (having any suitable structure) that is stored in a memory of the computer used to implement the content display system 203. The database can also store other information associated with each set of content data 300. This information is discussed in more detail below in the discussion of package files which can be used to convey such information from the content providing systems 203 to the content display systems 203. The package file editor mentioned above can be provided to each content providing system 203 to enable the content provider to easily create a package file for each set of content data 300 provided by that content provider.

Each package file includes a reference to the set of content data 300 (e.g., a network address) to which that package file corresponds. As mentioned above, each package file can also include a variety of other information. For example, the package file can include a specification of the format of the content data 300 (i.e., an indication of the types of content display systems 203 with which the set of content data 300 is compatible) and the type of the content data (e.g., an identification of a particular graphical format, as discussed above). (This information might be specified explicitly or implicitly; alternatively, this information may be passed to the content display system 203 separately from the package file.) The package file can additionally include a text description of the contents of the package file (this could be used, for example, in a user interface that lists descriptions of all of the sets of content data available to a content display system 203 or provided by a content providing system 203). The package file can also include information governing the presentation of the set of content data, such as screen position, special animation effects and display duration (the latter is shown by the View-Time attribute in the Example below). The update information (location and schedule) is also included in the package file. The package file can also include linking information (e.g., network address of an information source) used to implement a link option discussed in more detail below. The content data scheduling information discussed above can also be included in the package file. The package file can also include data structures that can be used to store auditing information, as discussed in more detail below. The package file can also include reference to one or more sets of control instructions 320, each set of control instructions 320 enabling display of the set of content data 300 by a content display system 203 having a particular architecture, or enabling display of clips of particular types.

22

The following Example illustrates how a package file for use with the invention could be constructed. The package file of this Example does not include all types of information that could be included in a package file; it is to be understood that other types of information (as discussed above, for example) could be included in such a package file, expressed in a similar manner to that shown in the Example. In this illustration, the package file is constructed in an object-oriented manner. Generally, each statement in the package file conforms to the following syntax:

keyword{attribute:value (1) . . . attribute:value(n)}

where "keyword" can be either PACKAGE or CLIP, "attribute" identifies one of the types of information discussed immediately above, and "value" is an identification of particular content for the type of information. There can be any number of "attribute:value" pairs in a statement. In the Example, each attribute:value pair is designated at right by a numeral enclosed in parentheses to aid in the description; this numeral does not form part of the package file shown in the Example.

Example

```

PACKAGE {
  Object-Id: 1 (1)
  Object-Type: 1 (2)
  Source: http://www.interval.com/~freiberg/
    Netscreen/Bookreviews/reviews.nss (3)
  Name: Book Reviews: Day 1 (4)
  Description: (5)
  Update-Frequency: 720 (6)
}
CLIP {
  Object-Id: 16919316 (7)
  Object-Type: 2 (8)
  Source: http://www.interval.com/~freiberg/
    Netscreen/Bookreviews/1%20Day%20Book/
    bookreview-1-a1.gif (9)
  Name: Anger (10)
  Description: Book Review (11)
  Update-Frequency: 0 (12)
  View-Time: 15 (13)
  Followup-URL: http://www.randomhouse.com/
    knopf/ (14)
  Linked-To-Following: 1 (15)
}
CLIP {
  Object-Id: 16919384 (16)
  Object-Type: 2 (17)
  Source: http://www.interval.com/~freiberg/
    Netscreen/Bookreviews/1%20Day%20Book/
    bookreview-1-a2.gif (18)
  Name: Emotional Intelligence (19)
  Description: (20)
  Update-Frequency: 0 (21)
  View-Time: 15 (22)
  Followup-URL: http://www.randomhouse.com/
    knopf/ (23)
  Linked-To-Following: 0 (24)
}

```

The first part of line 1 indicates that the following describes a package file. The remainder of line 1 and line 2 are used in debugging and are not relevant to the invention. Line 3 specifies a network address that identifies the location of the package file. (Note that the type of the package file is suggested in line 3 by the file extension .nss, though this extension is not necessary to specify the type. This extension can be used to implicitly specify the type of the package file to the content display system 203.) Lines 4 and 5 each give a description of the contents of the package file that can be used, for example, in a user interface to identify the package file. Line 6 specifies the frequency of acquisition of updates to the set of content data 300 and related instructions that are described by the package file. (In this Example, the frequency is specified in minutes.) The first part of line 7

6,034,652

23

indicates that the following describes a clip in the package file. The remainder of line 7 and line 8 are similar to lines 1 and 2. Line 9 specifies a network address that identifies the location of the clip. (Similarly to line 3, the type of the clip is suggested in line 9 by the file extension .gif, though, again, this extension is not necessary to specify the type.) Lines 10, 11 and 12 are similar to lines 4, 5 and 6. (Note that, in line 12, the specification of "0" for the update frequency indicates that the clip is never updated.) Line 13 specifies the duration of display for this clip. (In this Example, the duration is specified in seconds.) Line 14 specifies a network address of an information source to which a link can be established during display of this clip. (This aspect of the invention is described in more detail below with respect to the "more" option 603d of FIG. 6.) Line 15 specifies the number of additional clips that are part of this package file. Lines 16 through 24 are similar to lines 7 through 15.

When a content display system 203 receives a new package file from a content providing system 202, the content display system 203 first determines whether the contents of a corresponding package file (i.e., either the same package file or an earlier version of the package file) already exist as part of the database. This can be done by scanning a list of entries in the database that each indicate the presence of the contents of a particular package file. If the contents of a corresponding package file are not present, then a new entry is created in the list and the contents of the new package file are stored as part of the database (in accordance with the structure of the database). As part of the process of storing the contents of the package file, the contents are transformed into a form that is compatible with the architecture of the content display system 203 (this is enabled by the installation instructions 313 discussed above).

As described above, the database of content data and related information is constructed from a package file that can have a particular format, as illustrated in the Example above. However, generally, such a database can be constructed from files having any format (e.g., an ASCII file) that enables specification of the information described above that a package file includes.

Returning to FIG. 4, as indicated above, when a set of content data 300 is obtained (step 406), corresponding control instructions 320 and content data acquisition instructions 330 are also obtained (step 407) if such instructions have not already been acquired by the content display system 203. In particular, content data update instructions 333 can be obtained, so that updates to the set of content data 300 and/or the associated control instructions 320 and content data acquisition instructions 330 can be obtained in the future. As mentioned above, the content data update instructions 333 include a description of the location of the content providing system 202 from which the updates can be obtained as well as a schedule of times at which such updates should be obtained.

In the step shown in the block 408 (referred to hereinafter as step 408), a determination is made as to whether it is time to update the set of content data 300. The update schedule discussed above is used for this purpose. As long as the schedule indicates that no update need be obtained, the method 400 continues executing the step 408, thereby continuously monitoring whether an update need be obtained. The monitoring of step 408 could be implemented, for example, by a procedure that monitors the content display system computer clock and indicates that an update should be obtained when the clock time is equal to a time in the update schedule.

The update schedule can be established according to any desired criteria. For example, preferably, though not

24

necessarily, the step 408 (and the steps 409 and 410 discussed below and, as necessary, the steps 403 through 407 discussed above) of the method 400 operates at any time that the computer (or computers) with which the content display system 203 is implemented is on, even when the attention manager is not operating. Thus, the update schedule could be established so that updates are obtained during the middle of the night, when charges for communication with content providing systems 202 are cheaper. Preferably, then, at least this part of the content display system 203 is implemented on a computer that is always on, so that such cheap communications time can be utilized for obtaining updates. This can be particularly feasible if the content display system 203 is implemented on a client-server network in which at least the content data update instructions 333 are executed by a server computer which remains on at all times. If, however, the computer on which the content data update instructions 333 are executed is turned off at a time when an update is scheduled to be retrieved, then the update can occur immediately after the next time that the computer is turned on.

This aspect of the content data acquisition instructions 333 can be implemented, for example, using a communications daemon that is part of the content data update instructions 333. When the content data update instructions 333 are acquired by a content display system 203, the daemon is inserted into a startup file that is executed at the beginning of operation of the operating system of the computer with which the content display system 203 is associated. The daemon causes a connection to be made to each location from which the content data update instructions 333 indicate that an update is to be acquired. For example, if the computer uses a Windows operating system, the daemon initiates a WinSock TCP/IP connection to enable connection to be made to the locations of the updated sets of content data 300.

Returning to FIG. 4, once it is determined that an update of the set of content data 300 should be obtained, then, in the step shown in the block 409 (referred to hereinafter as step 409), the location of the appropriate content providing system 202 is ascertained from the scheduling information, and that location is accessed.

In the step shown in the block 410 (referred to hereinafter as step 410), a determination is made as to whether an updated set of content data 300 is available on the content providing system 202. If an updated set of content data 300 is not available, then the step 408 begins executing again, continuing until the update schedule indicates that it is again time to check for an updated set of content data 300. If an updated set of content data 300 is available, then the method 400 returns to the step 403, and an updated set of content data 300 and, if necessary, related control instructions 320 and content data acquisition instructions 330 are provided to the content display system 203 (i.e., an appropriate package file is provided to the content display system 203). As discussed above, the content display system 203 compares the version of the package file contents stored in the database to the contents of the version of the package file being newly provided, and makes changes to the database as necessary.

FIGS. 5A and 5B together are a flow chart of a method 500 that implements an attention manager according to another embodiment of the invention. Like the method 100 (FIG. 1), the method 500 is performed by a content display system 203 according to the invention which can be implemented, for example, using a digital computer that includes a display device and that is programmed to perform the functions of the method 500, as described below. Below, the method 500 is described as implemented on such a

6,034,652

25

digital computer, though the method 300 could be implemented on other apparatus. Steps in the method 300 that are the same as steps in the method 100 are shown by like-numbered blocks. Generally, the method 300 differs from the method 100 in that the method 300 provides a number of control options that enable the user to effect particular types of control of the attention manager. While the method 300 and the associated description below illustrate several control options that can be used with an attention manager according to the invention, it is to be understood that an attention manager according to the invention could include any of a number of other options not shown in FIGS. 5A and 5B, or described specifically herein.

The attention manager according to this embodiment of the invention can include any suitable user interface to enable the user to specify a control option. FIG. 6 illustrates a computer display screen 600 including one embodiment of such a user interface. The screen 600 displays, in addition to an image generated from a set of content data 300, a dialog box 601 that includes a list of available control options 602a through 602e. The dialog box 601 can remain on the screen 600 during the entire time that the attention manager is operating. The available control options 602a through 602e shown in the dialog box 601—as well as additional control options that could be, but are not, included in the dialog box 601—are discussed in more detail below.

The manner of selecting an option depends upon the available user input device(s). For example, a keyboard could be used to move a cursor to a desired option, which is then selected using the Enter key. Or, a mouse could be used to move a cursor to a desired option, then clicked to select the option. Or, a touch pen could be used to contact the screen 600 (if the screen 600 is a touch-sensitive screen) at an appropriate location to cause a desired option to be selected. Or, an audio command could be issued to a voice recognition system which causes the desired option to be selected.

One control option that can be used with an attention manager according to the invention enables the user to directly terminate operation of the attention manager. In FIG. 6, this is shown as the "exit" option 602a. In the method 300, this option is implemented using the step 107. As discussed above, selection of the "exit" option 602a causes the primary user interaction to begin again (block 101).

Another control option that can be used with an attention manager according to the invention enables the user to terminate display of the currently displayed set of content data and begin display of the next scheduled set of content data. In FIG. 6, this is shown as the "next" option 602b. In the method 300, this option is implemented by the step shown in the block 304.

Yet another control option that can be used with an attention manager according to the invention enables the user to terminate display of the currently displayed set of content data and begin display of the set of content data displayed immediately prior to the terminated set. In FIG. 6, this is shown as the "back" option 602c. In the method 300, this option is implemented by the steps shown in the blocks 303 and 311.

Still another control option that can be used with an attention manager according to the invention enables the user to terminate display of the currently displayed set of content data and remove that set of content data from the schedule so that the set will not be displayed in the future. This option is not shown in FIG. 6. In the method 300, this option is implemented by the steps shown in the blocks 303

26

and 312. In a particular embodiment, this option can be implemented so that the set of content data is precluded from being displayed only during the current operation of the attention manager. In another particular embodiment, this option can be implemented so that the set of content data is removed from the content display system 203 entirely, i.e., the set of content data is no longer available for display. In this embodiment, the set of content data could only become available for display again if the user takes affirmative steps to re-obtain the set of content data, as described above with respect to step 401 of the method 400 (FIG. 4).

Another control option that can be used with an attention manager according to the invention enables the user to prevent future display of the currently displayed set of content data until that set of content data has been updated. This option is not shown in FIG. 6. In the method 300, this option is implemented by the steps shown in the blocks 304, 313, 321, 322 and 323 (referred to hereinafter as steps 304, 313, 321, 322 and 323, respectively). If this option is selected in step 304, then an update flag is activated. The update flag can be a designated field associated with a particular set of content data in the database that contains all of the available sets of content data. As shown by step 321, the method 300 identifies, before display of a next set of content data in the schedule, the identity of that next set, and determines (step 322) whether the update flag has been activated for that set of content data. If the update flag has not been activated, then, in step 103, the set of content data is displayed. However, if the update flag has been activated, then, in step 323, a determination is made as to whether the set of content data has been updated since the last time that the set of content data was displayed. This step can be accomplished by checking an update monitor flag that can be a designated field of the database that is associated with the set of content data. If the update monitor flag indicates that the set of content data has been updated since the last time that the set of content data was displayed, then the set of content data is displayed (step 103). Otherwise, the method 300 returns to the step 321 to identify the next set of content data in the schedule.

Yet another control option that can be used with an attention manager according to the invention enables the user to specify a level of satisfaction with the currently displayed content data. This option is not shown in FIG. 6. In the method 300, this option is implemented by the steps shown in the blocks 303, 314 and 315. Depending upon the level of satisfaction indicated in the step 314, the schedule can be modified (step 315) to show the set of content data more, less or at different times than was previously the case. This option can be implemented in any appropriate manner; one way is described immediately below.

The content display system scheduling instructions 312 can include instructions that evaluate a probability function each time that a set of content data in the schedule is presented for display, and either display or not display the set of content data dependent upon the evaluation of the probability function. The probability function can include consideration of a variety of factors (e.g., the amount of time that has passed since a particular set of content data has been updated), but for implementation of the instant option, the probability function includes a term n^p , where n is a constant between 1 and 2, and p is a variable that represents a user's preference for a particular set of content data. Initially, the value of p is 0. Each time that a user indicates a like or dislike for a set of content data (by, for example, selecting an appropriate option in a dialog box such as the dialog box 601), the variable p is incremented or decremented,

6,034,652

27

respectively, by a predetermined amount. The content display system scheduling instructions 312 evaluate a stochastic probability function (e.g., a Gaussian probability function) using the evaluated probability function as an argument. If the result of evaluation of the stochastic probability function is "true", then the set of content data is displayed; if "false", then the set of content data is not displayed. As can be seen, then, initially (i.e., when $p=0$), the user has expressed no like or dislike of a set of content data and the set of content data is displayed or not according to other criteria. Incrementing or decrementing p (i.e., expressing like or dislike for a set of content data) causes the term n^2 to increase or decrease exponentially, thereby increasing or decreasing the likelihood that the set of content data will be displayed.

Still another control option that can be used with an attention manager according to the invention enables the user to establish a link with another information source. In FIG. 6, this is shown as the "more" option 602d. (In "wallpaper" embodiments of the invention, this option can be implemented so that any time the user clicks a mouse—or presses the "Enter" key on a keyboard—when the cursor is within the wallpaper, the link is made to the other information source.) In the method 300, this option is implemented by the steps shown in the blocks 306, 314 and 317. Links can be established to any of a variety of information sources and types of information sources. Typically, the link will be made to an information source that provides information that is related to the content data which was being displayed when the link was established. Upon selection by the user of this control option, the information source is accessed and additional information retrieved for presentation to the user. A link can be made, for example, to any information source that is part of a network which can be accessed by the computer with which the attention manager is being used (though it is not necessary that the link be made through a network). For example, the attention manager can be implemented so that links can be established to locations on the World Wide Web using the appropriate URLs. Such links can be established using any of a variety of Web browser software programs, such as the Navigator software program made by Netscape Communications Corp. Links are enabled by appropriately specifying the location (e.g., a network address) of the information source. The location of an information source (or locations of information sources) can be specified by associating the location with the set of content data, for example, in a package file as described and illustrated above.

As shown in the method 300, the attention manager continues to operate during the time that the link is established and the link is established to an information source from which it is possible to return to the attention manager (see step 317). The presentation of the new information to the user can include an appropriate user interface mechanism that allows the user to request such a return to the attention manager. However, the capacity to return to operation of the attention manager may not always exist. In that event, the step 317 is not part of the method 300; rather, the method 300 terminates after the step 316 and the user operates in the environment of the information source from that point forward. Such termination of the attention manager will frequently be the case where the link is made via a network to an information source.

Another control option that can be used with an attention manager according to the invention enables the user to obtain an overview of all of the content data available for display by the attention manager. This option is not shown

28

in FIG. 6, nor is it implemented in the method 300 of FIGS. 5A and 5B. The overview could be presented textually, pictorially or aurally. The overview information can be obtained either via a link to another information location (e.g., the location of the application manager 201) as described above or from a memory associated with the content display system 203, the overview information having been communicated to the content display system 203 when a set of content data was obtained.

Still another control option that can be used with an attention manager according to the invention enables the user to maintain display of the currently displayed set of content data 300 until such display is terminated by the user. This option is not shown in FIG. 6, nor is it implemented in the method 300 of FIGS. 5A and 5B. Upon selection of this option, an appropriate user interface could be made to appear that allows the user to specify termination of the display. After termination of the display, the attention manager resumes normal operation, i.e., the next set of content data 300 is displayed.

The dialog box 601 also includes an additional option, the "cancel" option 602e. Selection of the "cancel" option 602e causes the dialog box 601 to be removed from the screen 600. The dialog box 601 can be made to reappear again using any appropriate technique. For example, the application instructions 310 can include appropriate instructions to cause the dialog box 601 to reappear when the user makes an input to the computer using an input device.

As discussed above (see FIG. 2), usage of the attention manager can be audited using audit instructions 340 (FIGS. 3A and 3C) that can be supplied by the application manager 201 to the content display systems 203, either directly or via the content providing systems 202. The audit instructions 340 can include instructions that cause a content display system 203 to record, as the attention manager is used, particular information (audit information) regarding use of the attention manager (or compute such information from other, more basic information recorded by the attention manager). The audit information can be stored by the content display system 203 in an appropriately structured database. The audit information can include, for example, the identity of each set of content data 300 displayed by the attention manager, the number of times that a set of content data 300 was displayed by the attention manager, the frequency (e.g., number of times per week) that a set of content data 300 was displayed by the attention manager, the times at which a set of content data 300 was displayed by the attention manager, a user-expressed satisfaction level for a particular set of content data 300, and the last set of content data 300 displayed to a user before the user either "passively" (i.e., by making an input to the computer with an input device) or "actively" (i.e., by selecting a control option) terminated operation of the attention manager (of interest, since the user presumably was viewing the display screen when such interaction occurred). The audit instructions can also include instructions that compile and/or analyze the audit information in a desired manner. The audit instructions 340 can also include instructions that cause audit information to be transmitted to a remote site (e.g., the application manager 201 or a content providing system 202). These instructions can include scheduling instructions that govern when the audit information is so communicated (e.g., after periodic time intervals), as well as instructions that identify the location (e.g., network address) of the remote site. The transfer of audit information can be accomplished, for example, using a conventional electronic mail mechanism, as known to those skilled in the art. The audit

6,034,652

29

instructions 340 can also include instructions that enable the content display system 203 to display audit information. Additionally, the audit instructions 340 can include instructions that enable the user to disable the audit function entirely, or that enable the user to prevent audit information from being transmitted to the application manager 201 and/or to content providing systems 202. These last instructions could also be accompanied by operating instructions that provide a control option or options to the user, in a manner similar to that described above with respect to FIGS. 5A, 5B and 6, that enable the user to select disablement of the audit function. The audit instructions 340 can also include instructions that cause the database of audit information to be erased at an appropriate time, such as after the audit information has been communicated to a remote site.

Auditing of use of the attention manager can be useful to both users of the attention manager and content providers for a variety of reasons. Such auditing can be used, for example, to illustrate to content providers the value of the attention manager as a tool for disseminating the content provider's information, by showing the content providers how many content data display systems 203 are displaying the content provider's content data. The auditing can also give content providers insight into the interests of computer users, enabling the content providers to better target the information that the content providers provide. The auditing can also indicate to a user the amount and types of the information that the user has been receiving.

Various embodiments of the invention have been described. The descriptions are intended to be illustrative, not limitative. Thus, it will be apparent to one skilled in the art that certain modifications may be made to the invention as described without departing from the scope of the claims set out below. For example, though it is contemplated that an attention manager according to the invention will typically be used to occupy the peripheral attention of a human computer user, generally the attention manager can be used to occupy the attention of any sentient being. For example, the attention manager may be useful in occupying the attention of domesticated animals such as dogs or cats, or providing training (i.e., audio that can be repeated) for a "talking" bird such as a parrot.

We claim:

1. A system for engaging the peripheral attention of a person in the vicinity of a display device of an apparatus, comprising:

a content display system associated with the display device, the content display system including means for receiving a set of content data and a set of instructions for enabling a display device to selectively display, in an unobtrusive manner that does not distract a user of the apparatus from a primary interaction with the apparatus, an image or images generated from a set of content data, the content display system further including means for using the display device to selectively display the image or images using the set of instructions;

a content providing system including means for providing a set of content data to the content display system; means for providing to the content display system a set of instructions for enabling a display device to selectively display an image or images generated from a set of content data;

first communication means for enabling communication between the means for providing and the content display system;

second communication means for enabling communication between the content providing system and the content display system; and

30

means for auditing the display of sets of content data by the content display system.

2. A system for engaging the peripheral attention of a person in the vicinity of a display device of an apparatus, comprising:

means for acquiring a set of content data from a content providing system;

means for selectively displaying on the display device, in an unobtrusive manner that does not distract a user of the apparatus from a primary interaction with the apparatus, an image or images generated from the set of content data; and

means for detecting an idle period of predetermined duration, wherein the means for selectively displaying displays the image or images automatically after detection of the idle period.

3. A system as in claim 2, further comprising means for detecting a predetermined user interaction with the apparatus subsequent to detection of the idle period, wherein occurrence of the predetermined user interaction causes the means for selectively displaying to stop displaying an image or images generated from a set of content data.

4. A system for engaging the peripheral attention of a person in the vicinity of a display device of an apparatus, comprising:

means for acquiring a set of content data from a content providing system;

means for selectively displaying on the display device, in an unobtrusive manner that does not distract a user of the apparatus from a primary interaction with the apparatus, an image or images generated from the set of content data;

means for displaying one or more control options with the display device while the means for selectively displaying is operating;

means for selecting a displayed control option; and means for controlling aspects of the operation of the system in accordance with a selected control option.

5. A system as in claim 4, wherein:

the control option enables the user to request termination of operation of the system; and the means for controlling terminates operation of the system.

6. A system as in claim 4, wherein:

the means for selectively displaying further comprises means for scheduling the display of an image or images generated from a set of content data;

the control option enables the user to request display of a next image or images generated from a next set of content data; and

the means for controlling displays the next image.

7. A system as in claim 4, wherein:

the means for selectively displaying further comprises means for scheduling the display of an image or images generated from a set of content data;

the control option enables the user to request display of a previous image generated from a previous set of content data; and

the means for controlling displays the previous image.

8. A system as in claim 4, wherein:

a plurality of sets of content data are acquired by the system;

the means for selectively displaying further comprises means for scheduling the display of the image or images generated from the sets of content data;

6,034,652

31

the control option enables the user to remove a set of content data from the schedule; and
the means for controlling removes the set of content data from the schedule.

9. A system as in claim 4, wherein:

a plurality of sets of content data are acquired by the system, at least one of the sets of content data capable of being occasionally updated;

the means for selectively displaying further comprises means for scheduling the display of the image or images generated from the sets of content data;

the control option enables the user to prevent the display of an image generated from a designated set of content data until the designated set of content data has been updated; and

the means for controlling prevents the display of the image generated from the designated set of content data until the designated set of content data has been updated.

10. A system as in claim 4, wherein:

a plurality of sets of content data are acquired by the system;

the means for selectively displaying further comprises means for scheduling the display of the image or images generated from the sets of content data;

the control option enables the user to specify a satisfaction level for a currently displayed image from a current set of content data; and

the means for controlling revises the schedule in response to the specified satisfaction level.

11. A system as in claim 4, wherein:

the control option enables the user to establish a link with an information location; and

the means for controlling establishes the link with the information location.

12. A method for engaging the peripheral attention of a person in the vicinity of a display device of an apparatus, comprising the steps of:

acquiring a set of content data from a content providing system;

detecting an idle period of predetermined duration; and
selectively displaying on the display device, in an unobtrusive manner that does not distract a user of the apparatus from a primary interaction with the apparatus, an image or images generated from the set of content data, wherein the step of selectively displaying further comprises the step of displaying the image or images automatically after detection of the idle period.

13. A computer readable medium encoded with one or more computer programs for enabling acquisition of a set of content data and display of an image or images generated from the set of content data on a display device during operation of an attention manager, comprising:

acquisition instructions for enabling acquisition of a set of content data from a specified information source;

user interface installation instructions for enabling provision of a user interface that allows a person to request the set of content data from the specified information source;

content data scheduling instructions for providing temporal constraints on the display of the image or images generated from the set of content data, the content data scheduling instructions further comprising duration instructions for enabling specification of the duration of

32

time that the image or images generated from a set of content data can be displayed, wherein the duration instructions specify a duration of time that is dependent upon the particular time at which the image or images generated from a set of content data are displayed after the attention manager begins operating; and

display instructions for enabling display of the image or images generated from the set of content data.

14. A computer readable medium encoded with one or more computer programs for enabling acquisition of a set of content data and display of an image or images generated from the set of content data on a display device during operation of an attention manager, comprising:

acquisition instructions for enabling acquisition of a set of content data from a specified information source;

user interface installation instructions for enabling provision of a user interface that allows a person to request the set of content data from the specified information source;

content data scheduling instructions for providing temporal constraints on the display of the image or images generated from the set of content data, the content data scheduling instructions further comprising duration instructions for enabling specification of the duration of time that the image or images generated from a set of content data can be displayed, wherein the duration instructions specify a duration of time that is dependent upon the number of previous times that the image or images have been displayed during a continuous operation of the attention manager; and

display instructions for enabling display of the image or images generated from the set of content data.

15. A computer readable medium encoded with one or more computer programs for enabling acquisition of a set of content data and display of an image or images generated from the set of content data on a display device during operation of an attention manager, comprising:

acquisition instructions for enabling acquisition of a set of content data from a specified information source;

user interface installation instructions for enabling provision of a user interface that allows a person to request the set of content data from the specified information source;

content data scheduling instructions for providing temporal constraints on the display of the image or images generated from the set of content data, wherein the content data scheduling instructions further comprise sequencing instructions that specify an order in which the images generated from a set of content data are displayed; and

display instructions for enabling display of the image or images generated from the set of content data.

16. A computer readable medium as in claim 15, wherein the sequencing instructions further specify the duration of the display of each image or images generated from each set of content data.

17. A computer readable medium encoded with one or more computer programs for enabling acquisition of a set of content data and display of an image or images generated from the set of content data on a display device during operation of an attention manager, comprising:

6,034,652

33

acquisition instructions for enabling acquisition of a set of content data from a specified information source;
 user interface installation instructions for enabling provision of a user interface that allows a person to request the set of content data from the specified information source;

content data scheduling instructions for providing temporal constraints on the display of the image or images generated from the set of content data, wherein the content data scheduling instructions further comprise saturation instructions that constrain the number of times that the image or images generated from a set of content data can be displayed; and

display instructions for enabling display of the image or images generated from the set of content data.

18. A computer readable medium, for use by a content display system, encoded with one or more computer programs for enabling acquisition of a set of content data and display of an image or images generated from the set of content data on a display device during operation of an attention manager, comprising:

acquisition instructions for enabling acquisition of a set of content data from a specified information source;
 user interface installation instructions for enabling provision of a user interface that allows a person to request the set of content data from the specified information source;

34

content data scheduling instructions for providing temporal constraints on the display of the image or images generated from the set of content data;

display instructions for enabling display of the image or images generated from the set of content data;

content data update instructions for enabling acquisition of an updated set of content data from an information source that corresponds to a previously acquired set of content data;

operating instructions for beginning, managing and terminating the display on the display device of an image generated from a set of content data;

content display system scheduling instructions for scheduling the display of the image or images on the display device;

installation instructions for installing the operating instructions and content display system scheduling instructions on the content display system; and

audit instructions for monitoring usage of the content display system to selectively display an image or images generated from a set of content data.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 6,034,652
DATED : March 7, 2000
INVENTOR(S) : Freiburger et al.

Page 1 of 3

It is certified that error appears in the above-identified patent and that said Letter Patent is hereby corrected as shown below:

Title page.

Delete "Item [75] Inventors: Paul A. Freiburger, San Mateo, Calif.; Golan Levin, Staten Island, N.Y.; David P. Reed, Atherton, Calif.; Marc E. Davis, San Francisco, Calif.; Neal A. Bhadkamkar; Philippe P. Piernot, both of Palo Alto, Calif.; Todd A. Agulnick, San Francisco, Calif.; Sally N. Rosenthal, Palo Alto, Calif.; Giles N. Goodhead, Los Angeles, Calif."

and substitute

-- Item [75] Inventors: Paul A. Freiburger, San Mateo, Calif.; Philippe P. Piernot, Palo Alto, Calif.; Giles N. Goodhead, Los Angeles, Calif.; Neal A. Bhadkamkar, Palo Alto, Calif.; Todd A. Agulnick, San Francisco, Calif.; David P. Reed, Dover, Mass.; Golan Levin; Marc E. Davis, both of San Francisco, Calif.; Sally N. Rosenthal, Palo Alto, Calif.; --.

Delete "5,573,643 11/1996 Jackson 395/200.48"

and substitute

-- 5,572,643 11/1996 Jackson 395/200.48 --.

Column 5.

Line 65, delete "5 to 5C" and substitute -- 5A, 5B and 5C --;

Column 11.

Line 67, after "5B", insert --, 5C --;

Column 22.

Line 29, delete "Update Frequency" and substitute -- Update-Frequency --;

Line 34, delete "bookreview-1-a1.gif" and substitute -- bookreview-1-a1.gif --;

Column 24.

Line 39, delete "5A and 5B" and substitute -- 5A, 5B and 5C --;

Column 25.

Line 12, delete "5A and" and substitute -- 5A, 5B and 5C --;

Line 13, delete "5B";

**UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION**

PATENT NO. : 6,034,632
DATED : March 7, 2000
INVENTOR(S) : Freiburger et al.

Page 2 of 3

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 28.

Line 2, delete "5A and 5B", substitute -- 5A, 5B and 5C --;

Line 15, delete "5A and 5B", substitute -- 5A, 5B and 5C --;

Column 29.

Line 11, after "5B", insert --, 5C --;

Column 34.

Add the following claims:

19. A system as in Claim 2, further comprising means for non-volatily storing the set of content data.

20. A system as in Claim 2, further comprising:
means for indicating a time at which the means for selectively displaying is to begin display of the image or images; and
means for activating the means for acquiring at the indicated time, such that the means for selectively displaying displays the image or images in real time as the set of content data is acquired by the means for acquiring.

21. A system as in Claim 2, wherein the means for selectively displaying further comprises means for scheduling the display of the image or images generated from a set of content data.

22. A system as in Claim 2, further comprising means for updating the set of content data.

23. A system as in Claim 22, wherein the means for updating operates without disrupting use of the apparatus by the user during the time that the means for updating is operating.

24. A system as in Claim 22, wherein the means for updating obtains the updated set of content data from the content providing system.

25. A system as in Claim 24, wherein the means for updating operates automatically, without intervention by the user.

26. A system as in Claim 24, wherein the means for updating further comprises:
means for specifying the location of the content providing system; and
means for specifying the time at which an updated set of content data is to be obtained from the content providing system.

**UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION**

PATENT NO. : 6,834,652
DATED : March 7, 2000
INVENTOR(S) : Preilberger et al.

Page 3 of 3

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

27. A system as in Claim 2, wherein the means for detecting further comprises a location in a memory that stores data representing the duration of time since a last user interaction with the apparatus, the memory location being predetermined according to a method for operating the apparatus, such that an idle period is detected when the stored duration of time is equal to the predetermined duration.

28. A system as in Claim 2, further comprising:
means for interrupting a process being implemented by the apparatus at the time that the means for selectively displaying begins operating;
means for storing information representing the state of the process at the time of interruption; and
means for beginning operation of the process, using the stored state of the process, after the means for selectively displaying stops operating.

29. A system as in Claim 2, wherein the content data is video data.

30. A system as in Claim 29, wherein display of the video data produces a still image.

31. A system as in Claim 29, wherein display of the video data produces a moving image.

32. A system as in Claim 2, wherein the content data is audio data.

33. A system as in Claim 2, wherein the apparatus is a computer.

Signed and Sealed this

Thirteenth Day of November, 2001

Attest:

Nicholas P. Goenke

Acting Officer

NICHOLAS P. GOENKE
Acting Director of the United States Patent and Trademark Office

EXHIBIT 4

US006788314B1

(12) **United States Patent**
Freiberger et al.

(10) **Patent No.:** **US 6,788,314 B1**
 (45) **Date of Patent:** ***Sep. 7, 2004**

(54) **ATTENTION MANAGER FOR OCCUPYING THE PERIPHERAL ATTENTION OF A PERSON IN THE VICINITY OF A DISPLAY DEVICE**

4,845,658	7/1989	Gifford	364/900
5,105,184	4/1992	Pirani et al.	340/721
5,305,195	4/1994	Murphy	364/401
5,347,632	9/1994	Filepp et al.	395/200
5,436,637	7/1995	Gayraud et al.	345/705

(75) **Inventors:** **Paul A. Freiberger, San Mateo, CA (US); Golan Levin, San Francisco, CA (US); David P. Reed, Dover, MA (US); Marc E. Davis, San Francisco, CA (US); Neal A. Shusterman, Palo Alto, CA (US); Philippe R. Flornet, Palo Alto, CA (US); Todd A. Aquilino, San Francisco, CA (US); Sally N. Rosenthal, Palo Alto, CA (US); Giles N. Goodhead, Los Angeles, CA (US)**

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

WO	WO 93/19427	9/1993
WO	WO 96/30664	10/1996

OTHER PUBLICATIONS

Gomes, Lee, "Upstart's Internet 'TV' Has Microsoft Tuned In", Wall Street Journal, Aug. 1996.

Rigdon, Joan E., "Screen Savers Go Beyond Fish, Flying Toasters", Wall Street Journal, Feb. 13, 1996.

(73) **Assignee:** **Interval Research Corporation, Palo Alto, CA (US)**

(*) **Notice:** **Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.**

This patent is subject to a terminal disclaimer.

(21) **Appl. No.:** **09/838,803**

(22) **Filed:** **Mar. 20, 2000**

Related U.S. Application Data

(63) **Continuation of application No. 09/372,399, filed on Aug. 10, 1999, now abandoned, which is a continuation of application No. 08/620,641, filed on Mar. 22, 1996, now Pat. No. 6,034,652.**

(51) **Int. Cl.⁷** **G09G 5/12; G09G 5/14; G06F 15/16**

(52) **U.S. Cl.** **348/730; 709/218**


(58) **Field of Search** **345/2.1, 705, 710, 345/730; 707/10, 104, 501; 709/217, 218, 219**

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,868,675 2/1975 Firmin 340/334

16 Claims, 8 Drawing Sheets



US 6,788,314 B1

Page 2

U.S. PATENT DOCUMENTS

5,498,002	3/1996	Gechter	273/434	5,740,549	4/1998	Reilly et al.	705/14
5,572,643	11/1996	Judson	395/793	5,768,528	6/1998	Stamm	395/200.61
5,583,560	12/1996	Florin et al.	348/7	5,819,284	* 10/1998	Faber et al.	709/203
5,589,892	* 12/1996	Kaes et al.	725/43	5,913,040	6/1999	Rakavy et al.	395/200.62
5,724,424	* 3/1998	Gifford	705/79	6,034,652	* 3/2000	Freiburger et al.	345/730

* cited by examiner

U.S. Patent

Sep. 7, 2004

Sheet 1 of 8

US 6,788,314 B1

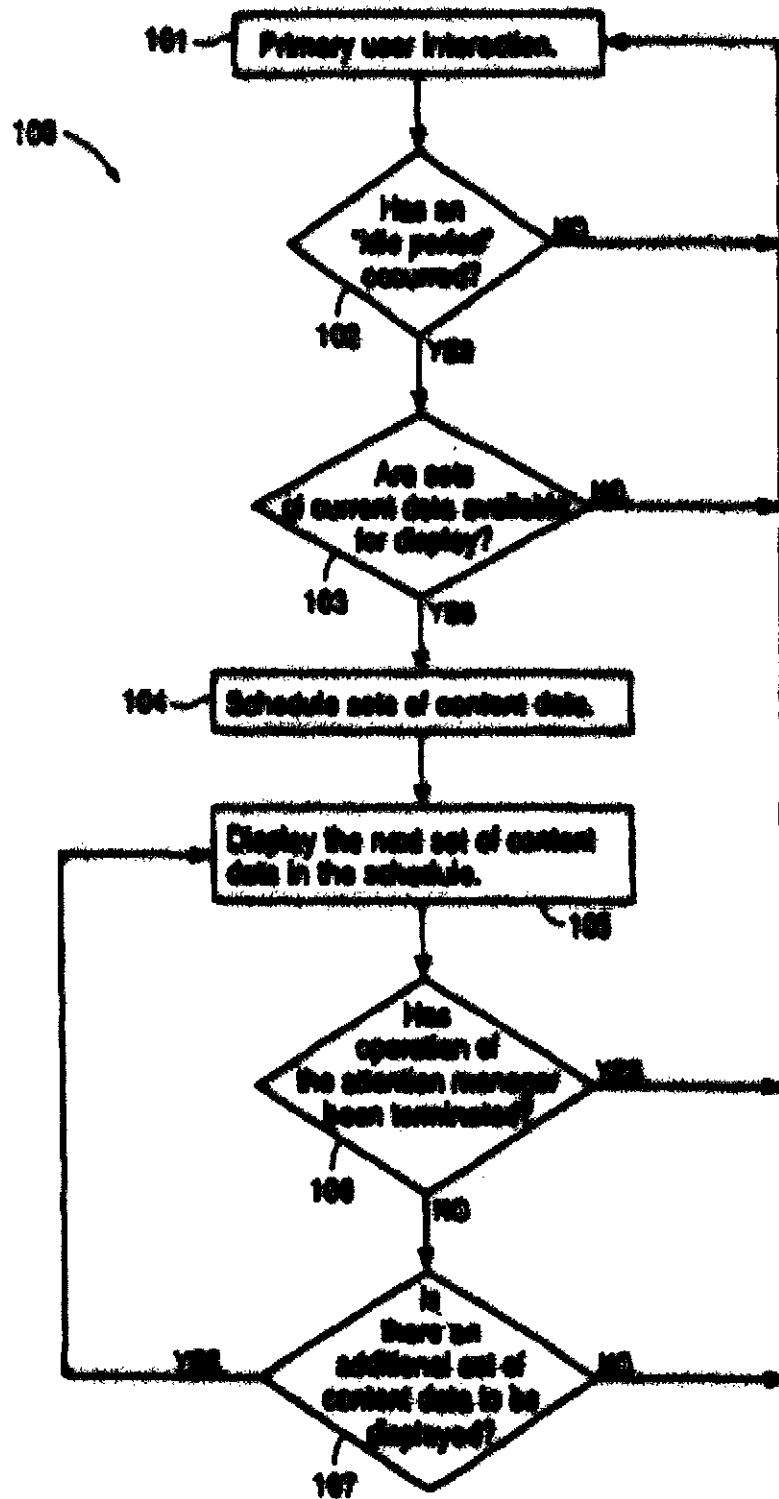


FIG. 1

U.S. Patent

Sep. 7, 2004

Sheet 2 of 8

US 6,702,314 B1




FIG. 2

U.S. Patent

Sep. 7, 2004

Sheet 3 of 8

US 6,798,314 B1




FIG. 3A

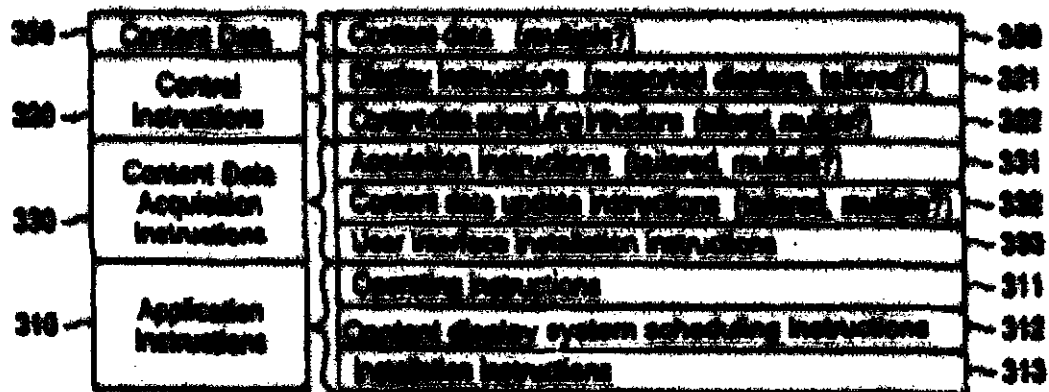


FIG. 3B




FIG. 3C

U.S. Patent

Sep. 7, 2004

Sheet 4 of 8

US 6,700,314 B1

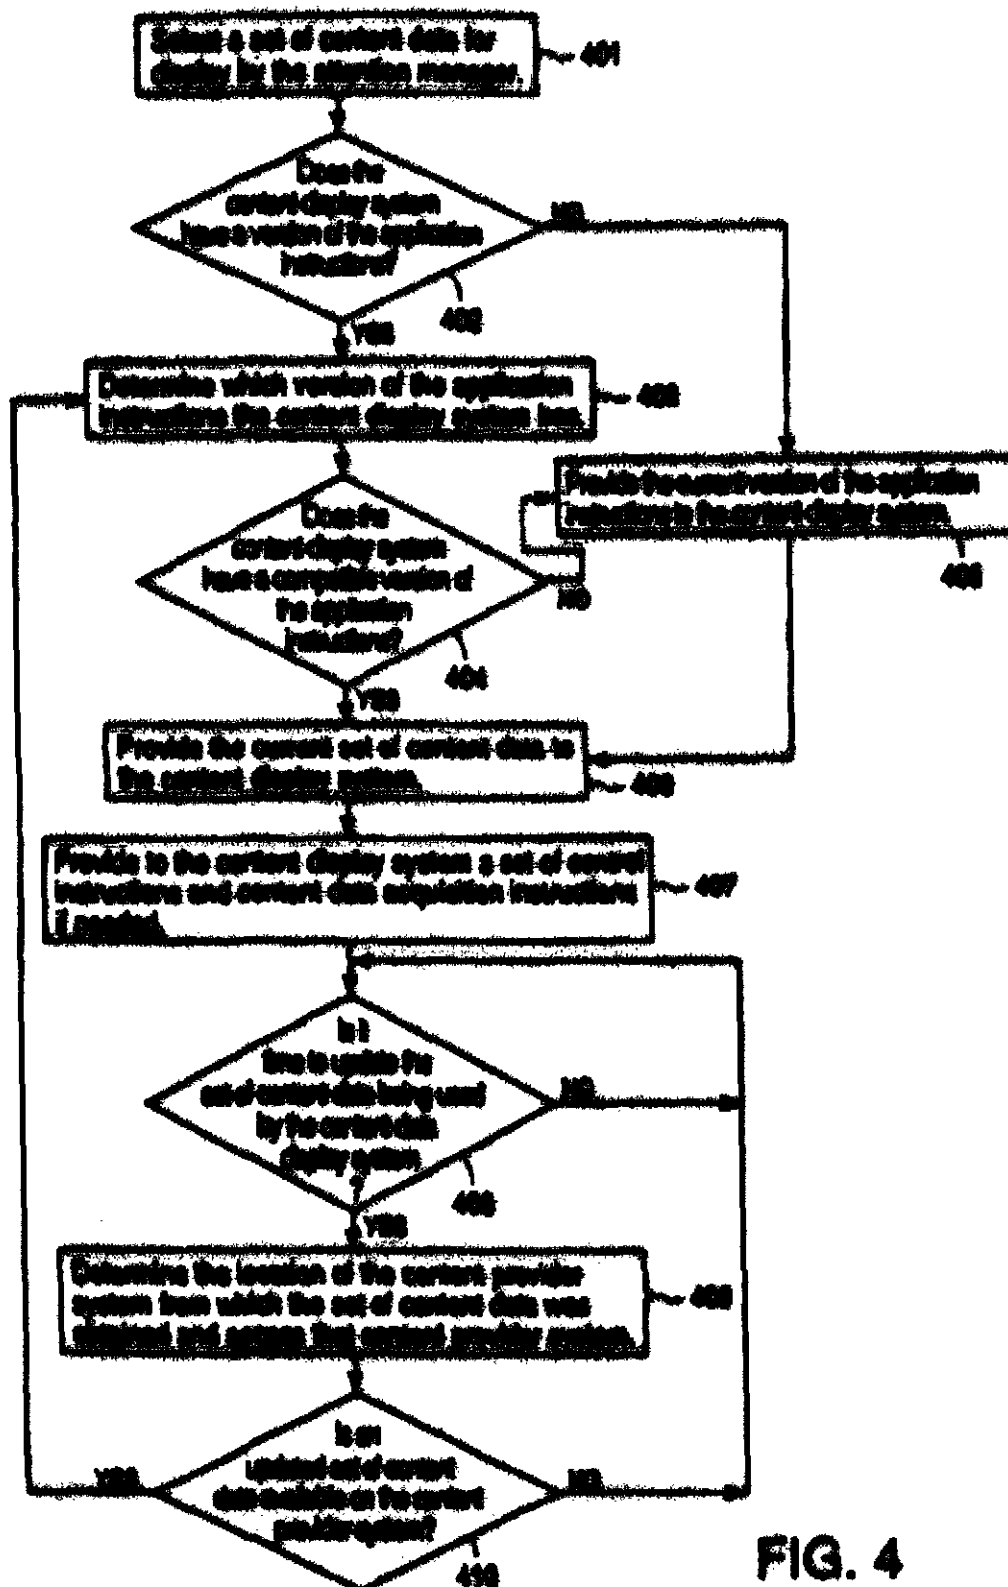


FIG. 4

U.S. Patent

Sep. 7, 2004

Sheet 5 of 8

US 6,788,314 B1


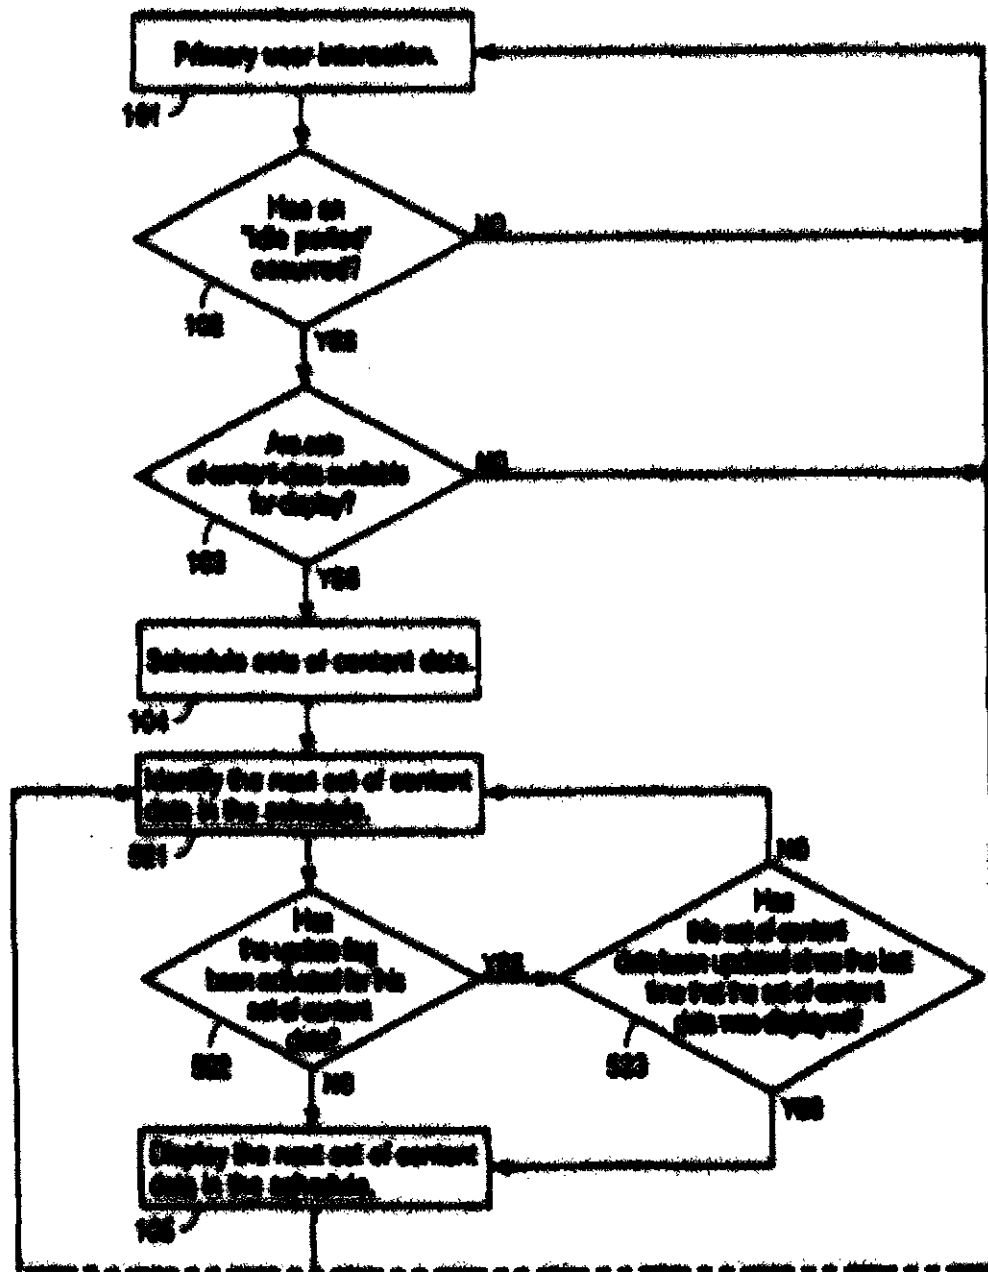


FIG. 5C



FIG. 5

FIG. 5A



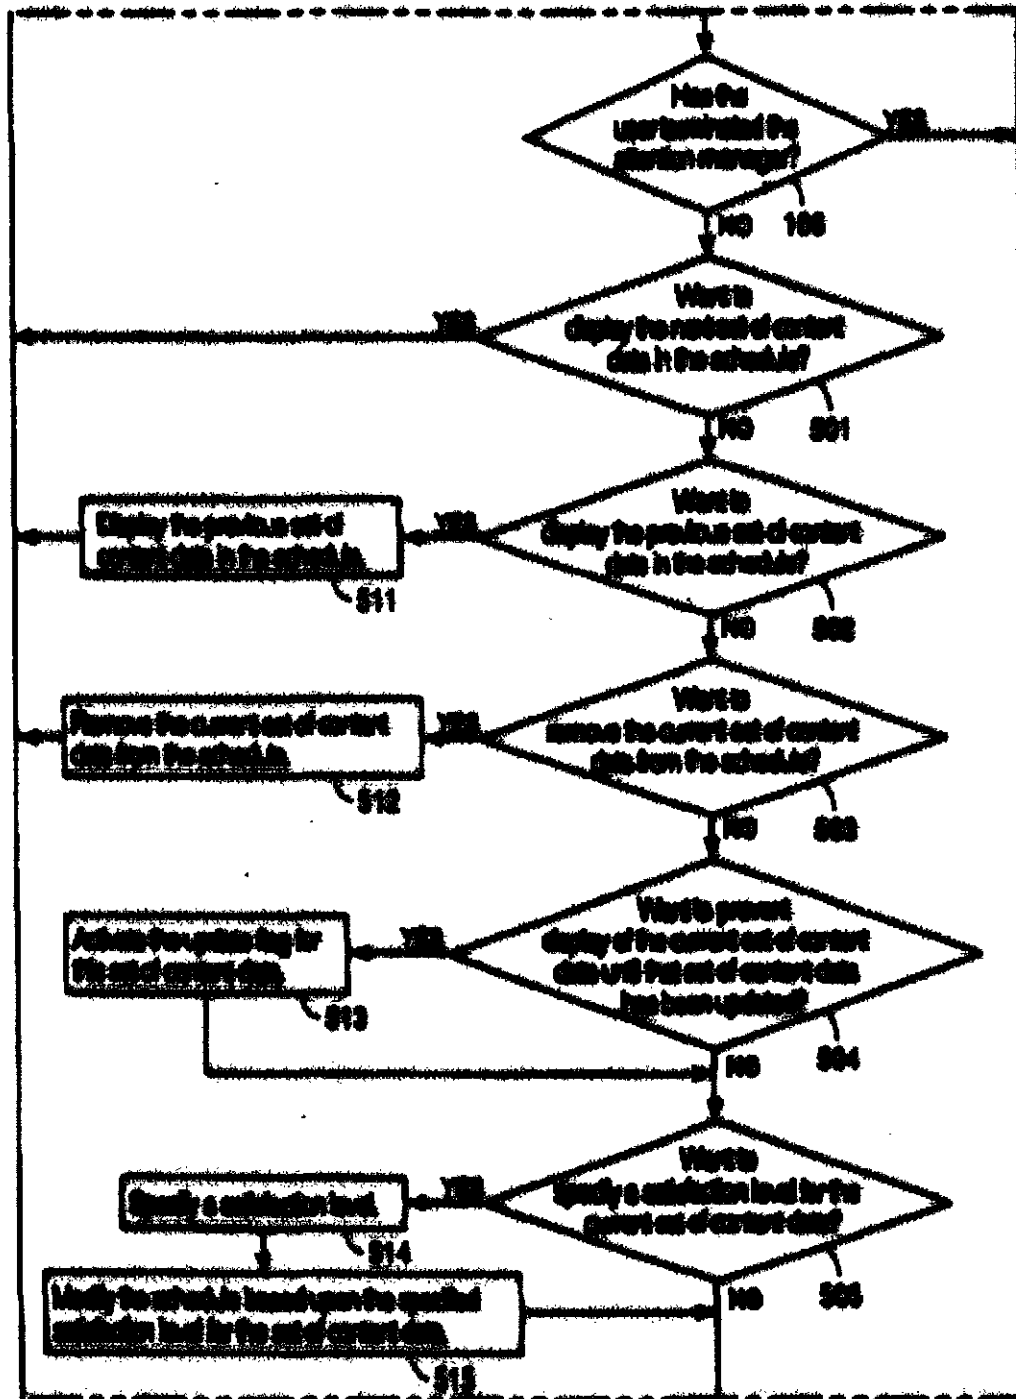
U.S. Patent

Sep. 7, 2004

Sheet 7 of 8

US 6,788,314 B1

FIG. 5B



U.S. Patent

Sep. 7, 2004

Sheet 8 of 8

US 6,780,314 B1




FIG. 6

US 6,788,314 B1

1

ATTENTION MANAGER FOR OCCUPYING THE PERIPHERAL ATTENTION OF A PERSON IN THE VICINITY OF A DISPLAY DEVICE

This application is a continuation of U.S. patent application Ser. No. 09/372,399, entitled ATTENTION MANAGER FOR OCCUPYING THE PERIPHERAL ATTENTION OF A PERSON IN THE VICINITY OF A DISPLAY DEVICE, filed Aug. 10, 1999, now abandoned, which is a continuation of U.S. patent application Ser. No. 08/620,641, filed Mar. 22, 1996, now U.S. Pat. No. 6,034,652.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to the engagement of the peripheral attention of a person in the vicinity of a display device such as the display monitor of a computer.

2. Related Art

Information providers of all sorts have an interest in presenting their information to information consumers and, in particular, to information consumers who may, or do, have an interest in the particular information provided by the particular information provider. At the same time, information consumers have an interest in accessing a wide variety of information and, in particular, information in which the information consumer may, or does have an interest. Given the extent to which computers now permeate society, and particularly in view of the escalation of networking of those computers in various ways, there is increasing recognition of the capability of using computers, and, in particular, computers (and other devices) that are interconnected in a network, as an information dissemination tool that can satisfy the interests of both information providers and information consumers.

For example, information providers have used public computer networks (e.g., the Internet) and private computer networks (e.g., commercial online services such as America Online, Prodigy and CompuServe) to disseminate their information. This information can be displayed to a computer user having access to the network directly in response to a request from the user or indirectly (i.e., without request by the user) as a result of another action taken by the user. While these methods of information dissemination and acquisition can be effective, they do not exhaust the possibilities.

In a different vein, historically, computers have frequently included screen saving mechanisms ("screen savers") intended to prevent the phosphors of a computer display screen from burning out when the same image remains on the screen for a long period of time, such as might occur during a long period of inactivity while the computer is operating. As computer display screen technology has progressed, the use of screen savers to preserve the display screen has become increasingly unnecessary. However, the use of screen savers has continued—even proliferated—likely due to the aesthetic or entertainment value provided by the imagery of many screen savers. Further, the use of "wallpaper" (i.e., a pattern generated in the background portions on a computer display screen) in computer display screens has also arisen, largely one would suspect because of the aesthetic or entertainment value of the wallpaper imagery. While the use of screen savers and wallpaper with computer displays appeals to many users because of the imagery they present to the user, screen savers and wallpaper have not heretofore been used as a means to convey infor-

2

mation from information providers to computer users. Further, screen savers and wallpaper have previously been implemented as relatively simple, self-contained computer application programs that are not typically integrated with other application programs or other aspects of computer operation. In particular, screen saver and wallpaper application programs have not been constructed to enable retrieval of display content from a remote location via a computer network.

SUMMARY OF THE INVENTION

An attention manager according to the invention presents information to a person in the vicinity of a display device in a manner that engages the peripheral attention of the person. Often, the display device is part of a broader apparatus (e.g., the display device of a computer). Generally, the attention manager makes use of "unused capacity" of the display device. For example, the information can be presented to the person while the apparatus (e.g., computer) is operating, but during inactive periods (i.e., when a user is not engaged in an intensive interaction with the apparatus). Or, the information can be presented to the person during active periods (i.e., when a user is engaged in an intensive interaction with the apparatus), but in an unobtrusive manner that does not distract the user from the primary interaction with the apparatus (e.g., the information is presented in areas of a display screen that are not used by displayed information associated with the primary interaction with the apparatus).

The information is embodied as one or more sets of content data. The sets of content data represent sensory data; typically, the sensory data is either video or audio data. Each set of content data is formulated by a content provider and made available for use by an attention manager according to the invention. Each content providing system can provide more than one set of content data. The content providing systems provide user interface tools that enable a particular set of content data to be requested. Once one or more sets of content data has been acquired, a content display system integrates scheduling information for all sets of content data to produce a schedule according to which an image or images corresponding to the sets of content data are displayed on a display device associated with the content display system.

A set or sets of instructions for enabling a display device to selectively display an image or images generated from a set of content data are also made available for use by the content display systems. Typically, the instructions enable images generated from content data to be displayed automatically, without user intervention, in a predetermined manner, thereby enhancing the capability of the invention to occupy the user's peripheral attention. Further, the attention manager can be implemented so that the instructions are automatically acquired (or updated, if necessary) each time a user requests acquisition of a set of content data, thereby making acquisition of the instructions transparent to the user of the attention manager and thus increasing the ease of use for the user. The instructions can include application instructions, control instructions and content data acquisition instructions. The application instructions can include operating instructions for beginning, managing, and terminating operation of the attention manager on a content display system, content display system scheduling instructions for scheduling the display of content data on a content display system, and installation instructions for installing the operating instructions and content display system scheduling instructions on a content display system. The control instructions can include display instructions for enabling

US 6,788,314 B1

3

generation of images from the content data on a particular type of display device or from a particular type of content data, and content data scheduling instructions for enabling temporal control of the display of the images generated from a set or sets of content data. The content data acquisition instructions can include acquisition instructions for enabling the acquisition of a set of content data, content data update instructions for enabling update of a previously acquired set of content data, and user interface installation instructions for enabling provision of a user interface that allows a person to request a set of content data from a content providing system. Each of the application, control and content data acquisition instructions could be acquired from a content provider, or any one or all of the sets of instructions could be acquired from an application manager that provides generic sets of instructions that can be tailored as necessary or desirable by a content provider. Additionally, audit instructions can be made available that enable monitoring of usage of the attention manager.

According to one aspect of the invention, an attention manager engages the peripheral attention of a person in the vicinity of a display device of an apparatus by acquiring one or more sets of content data from a content providing system and selectively displaying on the display device, in an unobtrusive manner that does not distract a user of the apparatus from a primary interaction with the apparatus, an image or images generated from the set of content data. According to a further aspect of the invention, the selective display of the image or images begins automatically after detection of an idle period of predetermined duration (the "screen saver embodiment"). This aspect can be implemented, for example, using the screen saver API (application program interface) that is part of many operating systems. According to another further aspect of the invention, the selective display of an image or images occurs while the user is engaged in a primary interaction with the apparatus, which primary interaction can result in the display of an image or images in addition to the image or images generated from the set of content data (the "wallpaper embodiment"). If multitasking is allowed by the apparatus (e.g., by the computer operating system) with which the attention manager is used, the attention manager can be implemented so that, when operation of the attention manager is terminated, the user is returned to the state of the primary interaction that existed when operation of the attention manager began. The attention manager can also be implemented so that, during operation of the attention manager, the user is presented with a number of options regarding further use of the attention manager. In particular, one of the options can allow additional information to be obtained that is related to the set of content data for which an image is being displayed. Where the attention manager is implemented as part of a network, this option can enable information to be obtained from a remote information source via the network. Another option that can be implemented allows a user to specify a satisfaction level for a set of content data from which an image or images is being displayed, thereby affecting the frequency with which that set of content data is used by the attention manager in the future.

According to another aspect of the invention, an attention manager that engages the peripheral attention of a person in the vicinity of a display device includes a content display system associated with the display device, a mechanism that can communicate with the content display system via a first communications mechanism to provide to the content display system a set of instructions for enabling the display

4

device to selectively display content data, and a content providing system that can communicate with the content display device via a second communications mechanism to provide a set of content data to the content display system. The content display system uses the provided set of instructions to selectively display on the display device an image or images generated from the provided content data. The attention manager according to this aspect of the invention can further include an application management system that can communicate via a third communications mechanism to provide to either the content providing system or the content display system one or more sets of instructions for enabling a display device to selectively display an image or images generated from a set of content data. In the former case, the content providing system can, in turn, communicate with the content display system to provide the one or more sets of instructions. The attention manager according to this aspect of the invention can be implemented, for example, using existing computer networks of information sources, such as the Internet (in particular, the World Wide Web) or commercial online services, advantageously making use of pre-existing hardware and software for enabling communication over those networks. Typically, though not necessarily, an attention manager according to this aspect of the invention will include multiple content display systems and multiple content providing systems. The content providing systems will each be capable of providing one or more sets of content data, so that, overall, there will be multiple available sets of content data which can be of different types. There can also be multiple sets of instructions for enabling a display device to selectively display an image or images generated from a set of content data, which sets of instructions may be tailored to display images from particular types of content data or to display content data using a particular display device.

According to yet another aspect of the invention, a computer readable medium can be encoded with one or more computer programs for enabling acquisition of a set of content data and display of an image or images generated from the set of content data on a display device during operation of an attention manager. The instructions of the computer program can include: i) acquisition instructions for enabling acquisition of a set of content data from a specified information source, ii) user interface installation instructions for enabling provision of a user interface that allows a person to request the set of content data from the specified information source, iii) content data scheduling instructions for providing temporal constraints on the display of the image or images generated from the set of content data, and iv) display instructions for enabling display of the image or images generated from the set of content data. The computer readable medium can also further include content data update instructions for enabling acquisition of an updated set of content data from an information source that corresponds to a previously acquired set of content data, the content data update instructions specifying where and when to obtain the updated set of content data. The content data scheduling instructions can specify, for example, the duration of time that the image or images generated from a set of content data can be displayed, an order in which the images generated from a plurality of sets of content data are displayed, a time or times at which the image or images generated from a set of content data can or cannot be displayed, and/or constraint on the number of times that the image or images generated from a set of content data can be displayed. The display instructions can be tailored to enable display of the image or images generated from a set of content data on a display device of a particular type, or display of an image or images generated from a set of content data of a particular type.

US 6,788,314 B1

5

According to still another aspect of the invention, a computer readable medium can be encoded with one or more computer programs for enabling a content display system to selectively display on a display device, in an unobtrusive manner that does not distract a person from a primary interaction with an apparatus associated with the display device, an image generated from a set of content data. The instructions of the computer program can include: i) operating instructions for beginning, managing and terminating the selective display of the image on the display device, ii) content display system scheduling instructions for scheduling the display of the image on the display device, and iii) installation instructions for installing the operating instructions and content display system scheduling instructions on a content display system. The computer readable medium can also further include audit instructions for monitoring usage of the content display system to selectively display an image generated from a set of content data.

The attention manager according to the invention is a new and useful mechanism for providing information to users of the attention manager. The attention manager provides information in which a user has expressed an interest and, importantly, information that the user might not otherwise expend adequate energy to obtain. The user can tailor the information provided by interacting with specific information sources to indicate interest in particular information provided by a specific information source while the user is perusing other information from that information source (as opposed to giving a general indication of interest in information on a particular subject or of a particular kind, from which indication information that matches the indicated interest is automatically provided from various information sources). The user can also choose information from a wide variety of information sources; in particular, when the attention manager is implemented using a network (e.g., the Internet), the user can acquire information from a wide variety of remote information sources. Additionally, the information is presented to the user in a manner that uses portions of the user's attention capacity that may otherwise be unused or filled with extraneous information.

The attention manager according to the invention also provides a new and useful information dissemination tool to content providers. The attention manager affords an opportunity to content providers to disseminate their information to users that are interested in receiving such information, enabling the content providers to provide better directed information dissemination. Moreover, the attention manager provides access to the previously unused attention capacity of those interested users. Additionally, the attention manager allows content providers to tailor particular aspects of the attention manager as desired by the content provider, such as the acquisition of updated sets of the content provider's content data (e.g., the frequency of such updates), the display scheduling and manner of display of the content provider's content data, and the user interface that enables users to specify acquisition of the content provider's content data.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flow chart of a method that implements an attention manager according to an embodiment of the invention.

FIG. 2 is a block diagram of a system for implementing an attention manager according to an embodiment of the invention.

FIGS. 3A, 3B and 3C are schematic diagrams illustrating the functional components of an application manager, a

6

content providing system and a content display system, respectively, according to an embodiment of the invention.

FIG. 4 is a flow chart of a method, according to an embodiment of the invention, for acquiring and updating sets of content data.

FIGS. 5, 5A, 5B and 5C together are a flow chart of a method that implements an attention manager according to another embodiment of the invention.

FIG. 6 illustrates a computer display screen including a user interface, according to one embodiment of the invention, that can be used to enable a user to specify a control option.

DETAILED DESCRIPTION OF THE INVENTION

According to the invention, an attention manager presents information to a person in the vicinity of a display device (or devices) in a manner that engages at least the peripheral attention of the person. "Display device", as used herein, encompasses any device that presents sensory stimulus to the person and includes, for example, computer video display devices, televisions and audio speakers. Further, here, "in the vicinity of" means any location with respect to the display device from which the person can perceive the information being presented. For example, if the information is being presented in a visual form, then "in the vicinity of" means any location from which the person can see the information. Or, if the information is being presented in an aural form, then "in the vicinity of" means any location from which the person can hear the information.

Often, the display device is part of a broader apparatus that can be utilized by a user for a primary interaction that is unrelated to the attention manager. (However, the attention manager can also be used with a display device that is not part of a broader apparatus, the user engaging in a primary interaction with the display device.) For example, the display device can be part of a computer that can be used to implement any of a number of application programs (e.g., word processing programs, computer games, spreadsheets, etc.). The person whose attention is engaged by the attention manager can be the user or another person in the vicinity of the display device. In one embodiment of the invention, the information is presented by the attention manager while a primary interaction is ongoing, but during inactive periods (i.e., when the user is not engaged in an intensive interaction with the apparatus). In another embodiment of the invention, the information is presented by the attention manager during active periods (i.e., when the user is engaged in an intensive interaction with the apparatus), but in an unobtrusive manner that does not distract the user from the primary interaction (e.g., the information is presented in areas of a display screen that are not used by displayed information associated with the primary interaction). Generally, then, an attention manager according to the invention makes use of "unused capacity" of a display device, "unused capacity" being defined broadly to include, for example, the embodiments mentioned above, i.e., both temporal (e.g., the first-described embodiment above) and spatial (e.g., the second-described embodiment above) dimensions.

The information is embodied by one or more sets of content data. Each set of content data is formulated by a content provider and made available by a corresponding content providing system for use with the attention manager. Each content providing system can provide more than one set of content data. Moreover, each set of content data can include one or more "clips", each clip being a definable

US 6,788,314 B1

7

portion of the set of content data that is used to generate a particular "image." The term "image" is used broadly here to mean any sensory stimulus that is produced from the set of content data, including, for example, visual imagery (e.g., moving or still pictures, text, or numerical information) and audio imagery (i.e., sounds). The content providing systems can also provide user interface tools that allow a user of the attention manager to specify that they want to obtain a particular set of content data. Once obtained, one or more images generated from the clips of one or more sets of content data are displayed by a content display system. The content display system integrates scheduling information associated with the sets of content data to produce a schedule according to which the images corresponding to the sets of content data are displayed for a particular user of the attention manager.

A set or sets of instructions for enabling a display device to selectively display images generated from one or more sets of content data are also made available to users of the attention manager. The instructions include application instructions, control instructions and content data acquisition instructions. Typically, the instructions enable images generated from content data to be displayed automatically, without user intervention, in a predetermined manner, thereby enhancing the capability of the attention manager to occupy the user's peripheral attention. Different sets of instructions can be formulated, such that only images generated from sets of content data that are compatible with a particular set of instructions can be displayed using that set of instructions. Typically, an application manager establishes a standard set or sets of instructions which content providers can tailor to fit their needs or desires.

As indicated above, the sets of content data represent sensory data, i.e., data that can be used to generate images as defined above. Typically, the sensory data is either video or audio data. The kinds of content data that can be used with the attention manager are virtually limitless. For example, video data that might be used as content data includes data that can be used to generate advertisements of interest to the user, moving and still video images which can be real-time or pre-recorded (e.g., nature scenes, pictures of family members, MTV music segments, or video from a camera monitoring a specified location, such as ski slopes or a traffic intersection, for conditions at that location), financial data (e.g., stock ticker information) or news summaries. Audio data that might be used as content data includes data that can be used to generate, for example, music or news programs (e.g., radio talk shows).

The attention manager according to the invention is useful both to users of the attention manager and to content providers. For users, the attention manager provides information to a user in which the user has expressed an interest. In particular, the attention manager provides information to a user that the user might not otherwise expend adequate energy to obtain. Additionally, the information is presented to the user in a manner that uses portions of the user's attention capacity that may otherwise be filled with extraneous information. Further, a variety of information can be displayed (i.e., images can be generated from more than one set of content data), so that the user does not have to choose particular information to the exclusion of all other information.

For content providers, the attention manager affords an opportunity to disseminate information to users that are interested in receiving such information, thus enabling the content providers to provide better directed information dissemination. Moreover, the attention manager provides

8

access to the previously unused attention capacity of those interested users. Further, since information from more than one content provider can be displayed, content providers are more likely to have their information displayed, since their information is displayed in addition to, rather than instead of, the information of other content providers, thereby reducing the need to compete with other content providers for the attention of the user.

FIG. 1 is a flow chart of a method 100 that implements an attention manager according to an embodiment of the invention. The method 100 is performed by a content display system according to the invention. The content display system can be implemented, for example, using a digital computer that includes a display device and that is programmed to perform the functions of the method 100, as described below. Below, the method 100 is described as implemented on such a digital computer, though the method 100 could be implemented on other apparatus.

As shown by block 101, initially (i.e., before operation of the attention manager begins), a user is engaged in a primary user interaction, e.g., a primary user interaction with a computer. Though shown in FIG. 1, the primary user interaction of block 101 does not form part of the method 100 according to the invention. "Primary user interaction" is to be construed broadly and, generally, includes any operation of the computer (or other apparatus with which the user is engaging in an interaction) other than operation that is part of the attention manager according to the invention. When the user is interacting with a computer, the primary user interaction includes any operation of the computer that occurs to enable or to support the performance of the function or functions that provide the basis for the user's use of the computer. For example, the primary user interaction can be the use of any of a variety of conventional application programs (e.g., word processing programs, spreadsheet programs, personal finance programs, game programs, drawing programs, online services and Web browsers, among others). The primary user interaction can also be, for example, simply the operation of a conventional computer operating system, such as the Windows (e.g., Windows 3.1, Windows NT or Windows 95) or DOS operating systems produced by Microsoft Corp. of Redmond, Wash. or the Macintosh operating system produced by Apple Computer, Inc. of Cupertino, Calif., among others. While, typically, the display device produces a display as a result of the primary user interaction, this need not necessarily be the case.

The method 100 actually begins with the block 102. In the step shown in the block 102 (referred to hereinafter as step 102), a determination is made as to whether an "idle period" has occurred. Generally, as used herein, "idle period" refers to a period of time of specified duration during which a specified condition does not occur. However, typically, the specified condition is one having the characteristic that failure of the condition to occur is indicative of an extended lack of intensive (or focused) interaction with the computer by the user ("user inactivity"). For example, the specified condition could be the lack of an input from an input device of the computer, e.g., the absence of striking a key on a keyboard, clicking a mouse, pressing on a touch-sensitive area of a touchscreen or issuing a voice command. Alternatively, the attention manager could be implemented with an apparatus that can monitor the environment of the apparatus (e.g., with a video camera) and evaluate the environment to ascertain that an "idle condition" (e.g., the viewing direction of the user of the apparatus is turned away from the apparatus by a specified amount for a specified period of time) has occurred, such idle condition triggering operation of the attention manager.

US 6,700,314 B1

9

Theoretically, any duration of time can be specified to define the idle period. However, practically, the duration of time necessary to constitute an idle period cannot be so short that the attention manager begins operating at times that inhibit the user's primary interaction with the computer or that distract or annoy the user. Further, the duration of time chosen, as indicated above, should be sufficiently long to indicate an extended lack of interaction with the computer, suggesting that the user is not engaged in an interaction with the computer that the user would not want to have interrupted. However, the duration of time should not be so long that, for periods of user inactivity of a typical duration, the amount of time that the attention manager operates is undesirably short. In sum, choosing the duration of time that defines an idle period involves a balancing of the above considerations. Illustratively, the idle period can be defined as a period of between thirty seconds and two minutes during which the specified condition (e.g., user interaction with an input device) does not occur.

While detection of the idle period can be implemented in any suitable manner, one way in which such detection can be implemented is by monitoring an idle timer that is part of a screen saver API (application program interface) that is, in turn, part of an operating system used to operate the computer. Such screen saver APIs are commonly found in current operating systems such as the Windows or Macintosh operating systems discussed above. The idle timer could be monitored and a signal that an idle period has occurred generated when the magnitude of the idle time as indicated by the idle timer reaches a predefined threshold.

Detection of an idle period as the basis for beginning operation of the attention manager is an indirect activation of the attention manager. In an alternative embodiment, step 102 of the method 100 is modified so that the attention manager is activated directly by the user. In other words, step 102 would consist of waiting for explicit direction from the user to begin operation of the attention manager. Such explicit direction could be enabled with an appropriate user interface, such as an on-screen icon or a menu selection, that is always present on the display screen of the display device as part of a standard interface that is provided by the operating system. Examples of such standard interfaces are the "Apple Menu" provided as part of the Macintosh operating system, and the "Start Menu" or desktop icons provided as part of the Windows 95 operating system.

Returning to FIG. 1, if, in step 102, an idle period has not occurred, then the primary user interaction continues (block 101). The method 100 continues executing the step 102 at predefined time intervals (typically very short time intervals), thereby continually and frequently checking for the occurrence of an idle period.

If, in step 102, an idle period is detected, then, in the step shown in the block 103 (hereinafter referred to as step 103), a determination is made as to whether there are any sets of content data available for use in generating a display. (Hereinafter, reference is sometimes made to "displaying content data" or "displaying a set of content data"; it is to be understood that this means displaying images generated using the content data or set of content data.) Herein, "content data" refers to data that is used by the attention manager to generate displays (e.g., video images or sounds, or related sequences of video images or sounds). A "set of content data" refers to a related set of such data that is used to generate a particular display. A "clip" refers to a definable portion of a set of content data that is used to generate a particular image; a set of content data can include one or more clips and, therefore, can be used to generate one or

10

more images. The acquisition of content data by the content display system is described in more detail below. Here, it is sufficient to note that, over time, an attention manager can acquire any number of sets of content data that can be displayed by the content display system.

If, in step 103, no sets of content data are available for display, then the primary user interaction continues (block 101). The method 100 continues executing the steps 102 and 103 at predefined time intervals, continually checking for the occurrence of an idle period and the acquisition of at least one set of content data.

If, in step 103, at least one set of content data is available for display, then, in the step shown in the block 104 (hereinafter referred to as step 104), the available sets of content data are scheduled for display by the content display system. (Alternatively, in other embodiments of the invention, scheduling of the sets of content data can occur before the method 100 begins. Such scheduling might be implemented, for example, so that each time a new set of content data is received by the content display system, the schedule is revised to include the new set of content data.) Typically, when the content display system acquires a new (or updated) set of content data, scheduling information for that set of content data is also acquired. Taken together, the scheduling information for all of the sets of content data is used to determine a schedule for display of the sets of content data by the content display system. Generally, determining a display schedule involves specifying the order in which the sets of content data are to be displayed and the duration of time for which each set of content data is to be displayed. The determination of the display schedule can also accommodate (to the extent possible) any special scheduling parameters for particular sets of content data (e.g., restrictions specifying when a particular set of content data must be displayed or cannot be displayed), mediating any conflicts between the display requirements of particular sets of content data. Often, though not necessarily, once the order and duration of display are established, the sets of content data are repetitively displayed by cycling through the display schedule repeatedly until operation of the attention manager is terminated. However, even where such iteration through the display schedule occurs, the display schedule can also accommodate scheduling parameters that delete sets of content data from the display schedule during particular iterations, thereby, for example, controlling the frequency with which particular sets of content data are displayed. The display schedule can be stored in an appropriately structured database, as known by those skilled in the art, that is stored in a memory of the computer used to implement the content display system.

Any appropriate set of rules, that can, for example, be arranged in any appropriate hierarchical manner, can be used for establishing a display schedule and, in particular, mediating conflicts between conflicting scheduling parameters associated with different sets of content data. For example, one rule for mediating conflicts may give preference to displaying sets of content data so that the sets of content data are displayed inversely to the order in which they were obtained by the content display system. This rule might be further specified so that a set of content data that has never previously been displayed by the attention manager is displayed prior to display of a set of content data that has been previously displayed, even though an update of the previously displayed set of content data has been obtained at a later time than that at which the never displayed set of content data was obtained. Another rule for mediating conflicts might resolve a conflict between two sets of content

US 6,788,314 B1

11

data having scheduling parameters that specify display at the same sequential position in the display schedule by randomly selecting one of the sets of content data to be displayed first during each iteration through the display schedule. Still another rule for mediating conflicts might establish a hierarchy of kinds of content data, with sets of content data of kinds at the top of the hierarchy being given preference for display over those at the bottom. Yet another rule or set of rules for mediating conflicts may involve performing some sort of analysis of the characteristics of the sets of content data that have been obtained by a particular content display system to ascertain preferences indicated thereby, and giving preference to sets of content data that are evaluated to be relatively more preferred. Scheduling rules of this kind would typically be part of the scheduling parameters provided independent of the content providers (i.e., in content display system scheduling instructions, as discussed elsewhere herein and, in particular, with respect to FIGS. 3A through 3C below).

Other scheduling rules, not directed to mediating conflicts between sets of content data, can also be used in determining a schedule. For example, any set of content data that has been initially obtained before a certain time and/or that has been last updated before a certain time (i.e., a set of content data that is "stale") can be automatically precluded from being inserted into the display schedule. This exclusion could further be restricted to apply only to certain sets of content data or content data of certain kinds. Similarly, the frequency with which a particular set of content data appears in a display schedule can be based upon how stale the set of content data is. Scheduling rules of this kind would typically be part of the scheduling parameters provided by a content provider for a set of content data (i.e., in tailored content data scheduling instructions, as discussed elsewhere herein and, in particular, with respect to FIGS. 3A through 3C below).

The particular scheduling rules used may be influenced by the characteristics of a particular embodiment of the attention manager, such as the available kinds of content data or the characteristics of the potential users of the attention manager. The particular scheduling rules used may also be influenced by the need or desire to simplify implementation of the scheduling rules.

Returning to FIG. 1, once the sets of content data have been scheduled for display, then, in the step shown in the block 100 (hereinafter referred to as step 100), a set of content data is displayed. The content display system is provided with one or more sets of display instructions to enable display of the set or sets of content data on the display device (as discussed elsewhere herein and, in particular, with respect to FIGS. 3A through 3C below).

After a set of content data has been displayed, then, in the step shown in the block 106 (hereinafter referred to as step 106), a determination is made as to whether operation of the attention manager has been terminated. Generally, operation of the attention manager can be terminated either directly or indirectly. Indirect termination of operation of the attention manager can be effected by, for example, causing operation of the attention manager to terminate when the specified condition (the non-occurrence of which is used to signal an idle period) occurs. For example, the attention manager can be terminated if the user makes an input to the computer using an input device, e.g., strikes a key on a keyboard, clicks a mouse, presses on a touch-sensitive area of a touchscreen or issues a voice command. For indirect termination, it may be desirable to add a further step or steps to the method 100 that, upon an indication that indirect termination should occur (e.g., the occurrence of the speci-

12

fied condition), asks the user to confirm that termination of the attention manager is, in fact, desired, and, if so, terminates the attention manager upon appropriate specified user input. In contrast to indirect termination, direct termination of operation of the attention manager can be effected by, for example, causing operation of the attention manager to terminate when the user selects a control option that specifies such termination, as described in more detail below with respect to FIGS. 3A, 3B and 6.

If, in step 106, operation of the attention manager has been terminated, then the primary user interaction begins again (block 101). The method 100 then begins executing the step 102 again, checking for the occurrence of an idle period.

If, in step 106, operation of the attention manager has not been terminated, then, in the step shown in the block 107 (hereinafter referred to as step 107), a determination is made as to whether there is an additional set of content data to be displayed. Typically, in operation of an attention manager according to the invention, there will always be another set of content data to be displayed, since, as discussed above, the sets of content data in the display schedule are iteratively displayed until operation of the attention manager is terminated. However, this need not be the case. For example, a limit can be established on the number of times that each set of content data can be displayed, or on the total number of times that any set of content data is displayed.

If, in step 107, there are no additional sets of content data to be displayed, then the primary user interaction begins again (block 101). The method 100 then begins executing the step 102 again, checking for the occurrence of an idle period.

If, in step 107, there are additional sets of content data to be displayed, then the method 100 returns to the step 102 and displays a set of content data in accordance with the previously determined display schedule. Steps 102, 106 and 107 are continuously performed, resulting in the continuous display of sets of content data, until either the user terminates the attention manager (step 106) or there are no more sets of content data to be displayed (step 107).

In another embodiment of the invention, a step could be added to the method 100, either in place of or in addition to the step 107, or as part of the step 106, that causes operation of the attention manager to terminate after the attention manager has been operating for a specified period of time.

Further, in another embodiment of the invention, an appropriate step or steps could be added to the method 100 so that, at a specified time, such as after each iteration through the display schedule, the method 100 returns to the step 102 and re-determines the display schedule.

As described above, when the method 100 ends, the primary user interaction (block 101) begins again. Preferably, the primary user interaction begins again with the status existent at the time that the method 100 began. Thus, the primary user interaction must be held in abeyance while the method 100 is operating. This can be accomplished by implementing the method 100 (or any other embodiment of the attention manager) with a content display system that is implemented on a computer that operates with an operating system that allows "multi-tasking" (here used to mean either the suspension of one program while one or more other programs operate, or the execution of one program simultaneously with the execution of one or more other programs). The Windows and Macintosh operating systems (mentioned above), among others, are operating systems having this characteristic. Where the attention man-

US 6,788,314 B1

13

ager is implemented using a screen saver API that is part of the operating system, such multitasking occurs automatically as a characteristic of the screen saver API, i.e., when operation of the attention manager ends, the user is returned to the status of the primary interaction existent at the time that the attention manager started operating. In multitasking operating systems that do not include a screen saver API, this feature of the invention can be implemented by use of an appropriately programmed device driver, as known by those skilled in the art, that monitors user interaction, suspending and restarting the primary user interaction at the beginning and end of operation of the attention manager.

The method 100 (FIG. 1) described above is an embodiment of the invention in which the attention manager presents information to a person (which can be the user or another person) in the vicinity of the display device during inactive periods when a user is not engaged in an intensive interaction with the computer (as indicated by the step 102 which checks for the occurrence of an "idle period" before beginning operation of the attention manager). As indicated above, in other embodiments of the invention, the attention manager presents information to the person during active periods, but in an unobtrusive manner. In such embodiments, video content data could be presented, for example, as "wallpaper" on the display screen of a video display monitor. Audio content data according to these embodiments could be presented in the same way as for the embodiments of the method 100 described above. For implementation of such embodiments of the invention, the step 102 of the method 100 could be modified to be a determination as to whether the attention manager has been activated (typically this would require direct activation by the user). Alternatively, step 102 could be eliminated altogether and the attention manager could be implemented to operate at any time that the computer is operating and sets of content data are available for display (step 100). For these embodiments, it is, as above, necessary that the content display system be implemented on a computer operated by an operating system that allows multi-tasking as described above. In particular, simultaneous operation of programs must be allowed, since the attention manager operates while the primary user interaction is ongoing (note that the relationships between the block 101 and the method 100 shown in FIG. 1 are not present in these embodiments of the invention).

Though not confined to such use, the attention manager according to the invention is envisioned as having particular use as a system implemented on, and used by, a network of computers. In such an implementation, each content providing system is implemented on a content provider computer. (It is possible to have more than one content providing system on a content provider computer.) Content display systems are implemented on user computers. The content provider computers and user computers are integrated together into a network such that each user computer can communicate with one or more of the content provider computers. The content provider computers need not (and typically would not) communicate with each other. Likewise, the user computers need not (and typically would not) communicate with each other. Further, each user computer need not communicate with all, or even more than one, of the content provider computers. For example, an attention manager according to the invention could be implemented so as to make use of a network such as the Internet. In particular, the graphical attributes of the World Wide Web would be particularly useful in enabling the provision of user interfaces that allow users to access the attention manager while visiting network sites of content providing systems.

14

FIG. 2 is a block diagram of a system 200 for implementing an attention manager according to an embodiment of the invention. The system 200 includes an application manager 201, a multiplicity of content providing systems, shown as Content Providers 1 through n (content providing systems 202a, 202b and 202c are illustrated in FIG. 2), and a multiplicity of content display systems, shown as Users 1 through n (content display systems 203a, 203b and 203c are illustrated in FIG. 2). Hereinafter, the content providing systems and content display systems are referred to generally using the numerals 202 and 203, respectively. In FIG. 2, the solid lines indicate that communication must occur in the system 200 and the dashed lines indicate that communication may occur. However, note that, in another embodiment of the invention, the application manager 201 is not present, and communication between any particular content display system and particular content providing system need not necessarily occur.

The application manager 201, content providing systems 202 and content display systems 203 can be implemented using appropriately programmed digital computers. Generally, the computers can be any conventional digital computers including an input device (such as a keyboard, mouse or touch screen), an output device (such as a conventional computer display monitor and/or one or more audio speakers), a processing device (such as a conventional microprocessor), a memory (such as a hard disk and/or random access memory), additional conventional devices necessary to interconnect and enable communication between the above-listed devices, and communications devices (e.g., a modem) for enabling communication with other computers of the system. For example, the application manager 201 and content providing systems 202 can be implemented using conventional server computers, while the content display systems 203 can be implemented using conventional client computers. The application manager 201, content providing system 202 and content display systems 203 could also themselves each be implemented by a client-server network of computers. Communication between the computers can be accomplished using any appropriate communication transmission lines, such as conventional telephone lines, or high speed data transmission systems such as T1, T3 or ISDN. The communication can be managed using any appropriate conventional networking methods (e.g., computer programs and protocols) and apparatus, as known by those skilled in the art. In particular, as described further below, the computers are programmed to enable the content display systems 203 to communicate with the content providing systems 202 and application manager 201 even without direct action by the user. In addition to being programmed to enable networking, each of the computers is also appropriately programmed, as described above and below, to perform the functions of the application manager 201, content providing systems 202 and content display systems 203, as appropriate.

FIGS. 3A, 3B and 3C are schematic diagrams illustrating the functional components of the application manager 201, a content providing system 202 and a content display system 203, respectively, according to an embodiment of the invention. Each of the functional components are represented by a set of instructions and/or data. (In particular, each of the sets of instructions may include, if appropriate, data related to accomplishment of the functions associated with the set of instructions; similarly, a set of content data may include, if appropriate, instructions that enable generation of an image from the set of content data.) Each of these sets of instructions and/or data can be embodied in an appropriate com-

US 6,788,314 B1

15

puter program or set of computer instructions (the latter capable of including computer instructions and data), or an appropriate set of data configured for use by a set or sets of instructions (e.g., computer program) that must interact with the set of data in order to implement the attention manager.

The application manager 201 stores a variety of instructions for use in implementing the attention manager. As shown in FIG. 3A, generally, the application manager 201 stores application instructions 310, control instructions 320, and content data acquisition instructions 330 that can be disseminated to the content display systems 203 and content providing systems 202 as necessary or appropriate. The application manager 201 can also store audit instructions 340 that can be used to enable monitoring of usage of the attention manager.

The application instructions 310 include operating instructions 311 for beginning, managing and terminating operation of the attention manager on a content display system 203, and content display system scheduling instructions 312 for scheduling the display of content data on a content display system 203. The method 100 (FIG. 1) described above is one embodiment of such application instructions 310. The application instructions 310 also include installation instructions 313 that enable the other instructions used by the attention manager to be implemented using the hardware that is part of and associated with a particular content display system 203. The installation instructions 313 can be implemented as known by those skilled in the art. For example, the installation instructions 313 can be a "plug-in" or "helper" application program (such as a helper application that can be used with the Navigator and Mosaic software programs made by Netscape Communications Corp. of Mountain View, Calif.) that is used to process instructions or data of a particular type—in this case, instructions to implement the attention manager, and content data for use with the attention manager, that can be installed on the hardware of a particular content display system 203. There can be a multiplicity of such helper applications, each capable of operating on particular hardware that could be used to implement a content display system 203. The helper application enables the software program (i.e., Navigator or Mosaic) being used to access the sites of content providers to process references (e.g., Universal Resource Locators, or "URLs") to the particular type of instructions and/or data, so that sets of content data (including updated sets of content data) and the application instructions 310, control instructions 320 and content data acquisition instructions 330 (including updated versions of those instructions) can be acquired.

The control instructions 320 include display instructions 321 and content data scheduling instructions 322, as described in more detail below, that are typically enhanced by content providers in a particular manner that is appropriate for the content data that the content providers provide. The application manager 201 can (and typically does) store and disseminate multiple distinct sets of control instructions 320. Generally, the display instructions 321 of a particular set of control instructions 320 enable display of content data on a particular type of display device (e.g., a particular type of computer video display or a particular type of audio speaker) or display of a particular type of content data. Display instructions 321 that can be used with a particular display device are typically already developed by third parties (e.g., the maker of the display device) and are readily available. Tailoring of the display instructions 321 to display particular types of content data (such as instructions for displaying content data that is in the GIF format or the

16

format of AutoDesk Animator FLC files) can be done by either the application manager 201 or a content provider. The content data scheduling instructions 322 provide temporal constraints on the display of particular sets of content data. As stored by the application manager 201, the content data scheduling instructions 322 are usually the same for each set of control instructions 320 and provide a generic set of scheduling instructions that can be tailored by a content provider.

The content data acquisition instructions 330 include acquisition instructions 331 for enabling the initial acquisition of a set of content data and instructions for implementing the attention manager, and content data update instructions 332 for enabling update of previously obtained sets of content data and attention manager instructions. The acquisition instructions 331 and content data update instructions 332 are generic sets of instructions that can be tailored by a content provider. The content data acquisition instructions 330 can also include user interface installation instructions 333 that enable content providers to install a user interface in the information environment (e.g., Web page) of the content provider so that users can request sets of content data from the content provider. Such user interface installation instructions are conventional and readily available for use with the attention manager of the invention.

As shown in FIG. 3B, the content providing systems 202 store one or more sets of content data 300 that can be disseminated to content display systems 203 as requested. The content providing systems 202 can also store the application instructions 310, control instructions 320, and content data acquisition instructions 330 described above.

As indicated above, each set of content data 300 defines a related group of data that is used to generate a particular display and includes one or more clips that each represent a definable portion of the set of content data that is used to generate a particular image. The content data 300 represents sensory data and can be, for example, video or audio data. A particular set of content data 300 can be formulated in different versions that are each compatible with content display systems 203 having particular characteristics. In particular, the characteristics of the display device of a content display system 203 can affect the formulation of a set of content data 300. For example, for computer video display monitors, the formulation of a version of a set of content data 300 can depend on the size of the display screen (e.g., horizontal length by vertical length), the display resolution (e.g., the number of horizontal pixels by the number of vertical pixels), the color depth (number, e.g., 256, of possible colors) and the characteristics of the display drivers for the display device. The formulation of a version of a set of content data 300 could also depend upon the operating system being used by the computer on which the content display system 203 is implemented or other characteristics of the computer, such as the speed with which the display device can be operated (insofar as that speed is affected by the characteristics of the computer such as processor speed). Generally, a set of content data 300 can be formatted as known by those skilled in the art in view of the above considerations.

As indicated above, the control instructions 320 (as well as the content data acquisition instructions 330) are typically enhanced by content providers as appropriate for particular content data. The manner in which these instructions can be tailored by content providers is desirably required to conform to a specified format. Below, a description is given of package files that can be used for tailoring the control instructions 320 and content data acquisition instructions

US 6,788,314 B1

17

300, as well identifying the location of content data. These package files can be created using an appropriate computer program (package file editor) that can be provided by, for example, the application manager 201 and that enables this tailoring to be accomplished easily and according to the specified format.

The content provider can tailor the content data scheduling instructions 323 to indicate the duration of time that a particular set of content data can be displayed ("duration instructions"). Generally, the duration instructions can be arbitrarily complex and can vary in accordance with a variety of factors, including, for example, the particular time at which the set of content data 300 is displayed after the attention manager begins operating, or the number of previous times that the set of content data 300 has been displayed during a continuous operation of the attention manager. The content provider can also tailor the content data scheduling instructions 323 to indicate an order in which the clips of a set of content data 300 are displayed, as well as the duration of the display for each clip ("sequencing instructions"). The content provider can also tailor the content data scheduling instructions 323 to indicate particular times or ranges of times at which a set of content data 300 can or cannot be displayed ("timing instructions"). These times can be absolute (e.g., a particular clock time on a particular day, a particular day or days during a week, after or before a specified date) or relative (e.g., not before or after a specified duration of time since the attention manager began operation, first or not first among the sets of content data 300 to be displayed, not after a particular kind or set of content data 300). The content provider can also tailor the content data scheduling instructions 323 to specify a maximum number of times that the set of content data 300 can be displayed after the attention manager begins operating or a maximum number of times that the set of content data 300 can be displayed over any number of operations of the attention manager ("saturation instructions").

The content provider can also tailor the display instructions 324 to display a particular set or sets of content data. The display instructions 324 can be tailored, for example, according to the type or types of the content data. The type of content data indicates the manner in which an image or images are generated from the content data (i.e., how the bit patterns in a particular clip are transformed into an image). The type of content data is typically established as a consequence of the manner (e.g., with a particular software application program such as the Photoshop or Premiere programs produced by Adobe Systems of Mountain View, Calif.) in which a particular clip is created. The installation instructions 313, discussed above, enable content data of different types to be obtained by the attention manager. Generally, the possible types of content data can be confined to an enumerated set of standard data types, such as the Mime data types used with the World Wide Web. As will be more readily understood from the description below, the type of content data can be specified, for example, in a field of the clip part of a package file.

The ability to tailor sets of content data 300 and associated control instructions 330 for particular content display systems 200, before the sets of content data 300 and control instructions 330 are provided to those content display systems 200, is advantageous because it allows the tailoring to be done once, by the content providing system 202 or the application manager 201, rather than multiple times, once by each content display system 200 that uses the set of content data 300 and associated control instructions 330.

The content data acquisition instructions 330—in particular, the content data update instructions 333—are also

18

tailored by content providers as appropriate for particular sets of content data 300. In particular, the content provider can tailor the content data acquisition instructions 330 to indicate where and when to obtain an updated set of content data 300. For example, the indication of where to obtain an updated set of content data 300 can be accomplished by specifying an appropriate network address of a content providing system 202. The network addresses can be specified by, for example, a URL used to identify, for example, an HTML file, an applet (a short application program written in Java or other suitable programming language), a script based on CGI or other suitable mechanism, or any other resource (i.e., computer program or set of data). The indication of when to obtain an updated set of content data 300 can be accomplished by specifying a time or times, either absolute time or times (i.e., particular dates and times during the day) or relative time or times (e.g., one month after the last acquisition/update of the set of content data 300). For example, the update schedule could be established to obtain updates every hour, every day or every week. Or, the update schedule could be established to obtain updates upon the occurrence of a particular event, such as a specified percentage increase or decrease in a stock market index. In general, the particular update schedule used will depend upon the character of the content data with which the update schedule is associated, e.g., content data representing stock prices would probably be updated more frequently than content data representing an advertisement.

As shown in FIG. 3C, the content display systems 200 store the application instructions 310, control instructions 320, and content data acquisition instructions 330 described above. The application instructions 310 use the control instructions 320 to display sets of content data 300 that are obtained (and updated, if appropriate) by the content data acquisition instructions 330. The application instructions 310 and control instructions 320 are discussed generally, and with respect to particular embodiments, in more detail above, while an embodiment of the content data acquisition instructions 330 is described below.

FIG. 4 is a flow chart of a method 400 according to the invention for acquiring and updating sets of content data, i.e., the method 400 is an embodiment, at least in part, of the acquisition instructions 331 and content data update instructions 332 of the content data acquisition instructions 330 discussed above with respect to FIGS. 3A through 3C. In the method 400, the steps shown by blocks 403 through 407 can be implemented in the acquisition instructions 331 and the steps shown by blocks 408 through 410 can be implemented in the content data update instructions 332. Generally, the steps of the method 400 can be implemented on an appropriately programmed digital computer that is programmed to perform the functions of the method 400, as described below. Below, the method 400 is described as implemented on such a digital computer, though the method 400 is not limited to such an implementation. The method 400 necessitates communication between a content display system 200 and one or more content providing systems 202. As will be understood by those skilled in the art of digital computer programming for computer network communications, when the method 400 is implemented using a programmed digital computer, particular steps of the method 400 could be implemented on either a content display system 200 or a content providing system 202.

In the step shown in the block 401 (referred to hereinafter as step 401), a set of content data is selected for display by the attention manager. Initially, in step 401, particular sets of content data are obtained as a result of direct request by the

US 6,788,314 B1

19

user. Any appropriate user interface can be used for enabling a user to directly request a particular set of content data. For example, Web pages on the World Wide Web could include graphical buttons for enabling users that visit the Web page to request particular sets of content data. Selection of a button on a Web page results in an indication to the appropriate content providing system 203 that the requesting content display system 200 has requested the set of content data corresponding to the selected button to be transferred to the content display system 200. The user interface instructions 333 discussed above, that can be provided to each content providing system 203, can be used to create the user interface.

Selection of a set of content data in step 401 causes a set of acquisition instructions 331 to be transferred to the content display system 200. The acquisition instructions 331 include information identifying the site from which the set of content data can be obtained, as well as the site or sites from which instructions (e.g., application instructions 310, control instructions 320, content data acquisition instructions 330 and audit instructions 340) for implementing the attention manager can be obtained. These sites can be the same or different sites. The sites can be identified by, for example, using UMLs, as described above. The acquisition instructions 331 can also include instructions for establishing an appropriate user interface (e.g., a desktop icon) in the content display system 200 that enables a user to cause the installation instructions 313 to be executed, thereby installing the attention manager in the content display system 200.

In the step shown in the block 403 (referred to hereinafter as step 403), a determination is made as to whether the requesting content display system 200 has the application instructions 310 (FIGS. 3A through 3C) that enable operation of the attention manager and scheduling of sets of content data 300. If the content display system 200 does not have the application instructions 310, then, in the step shown in the block 406 (referred to hereinafter as step 406), the content display system 200 uses the appropriate site identification provided by the content providing system 203 to obtain a version of the application instructions 310 (typically the most current version of the application instructions 310 that is compatible with the set of content data 300 requested by the user). The application instructions 310 can be provided by the content providing system 203 from which the set of content data 300 is being obtained. Alternatively, the application instructions 310 can be provided directly to the content display system 200 by the application manager 201 (or from some site other than a content provider or the application manager 201) by causing an appropriate instruction to be issued to the application manager 201 (or other site) by either the content providing system 203 or the content display system 200.

If the content display system 200 does have the application instructions 310 (step 403), then, in the step shown in the block 403 (referred to hereinafter as step 403), a determination is made as to which version or versions of the application instructions 310 the content display system 200 has. As indicated elsewhere, a particular set of content data 300 can (and typically will) be updated from time to time, thereby creating different versions of the set of content data 300. Likewise, the application instructions 310 can also be updated, thereby creating different versions of the application instructions 310. In general, a set of content data 300 can be updated without regard to whether the set of content data is compatible with all versions of the application instructions 310 (though the set of content data 300 must be compatible with at least one version of the application

20

instructions 310). Likewise, the application instructions 310 can be updated without regard to whether any particular set of content data 300 is compatible with that version of the application instructions 310. Moreover, particular versions of the application instructions 310 may be compatible only with sets of content data 300 of certain types. Consequently, a particular content display system 200, even though the content display system 200 has the application instructions 310, may not have a version of the application instructions 310 that is compatible with the type and/or version of the set of content data 300 being requested.

It is necessary, therefore, to determine whether the content display system 200 has a version of the application instructions 310 that is compatible with the type and version of the set of content data 300 being requested so that, if necessary, a compatible set of application instructions 310 can be provided to the content display system 200. In the step shown in the block 404 (referred to hereinafter as step 404), this determination is made. If the content display system 200 does not have a compatible version of the application instructions 310, then, in step 406, the content providing system 203 (or, for example, the application manager 201) provides to the content display system 200 a version of the application instructions 310 (typically the most current version) that is compatible with the requested set of content data 300.

Alternatively, in step 404, a determination could be made as to whether the version of the application instructions 310 that the content display system 200 has is the most current version of a set of compatible application instructions 310. If the version is not the current version, then the content providing system 203 provides the current version (step 406), even if the version that the content display system 200 already had is compatible with the newly acquired set of content data 300.

Preferably, updated sets of application instructions 310 are made downwardly compatible with previous sets of application instructions 310, so that the updated application instructions 310 can be used with any previously obtained sets of content data that are compatible with a previous set of application instructions 310. If downward compatibility is not maintained, the updated set of application instructions 310 can replace a previous set of application instructions 310 and incompatible sets of content data can be removed from the schedule of sets of content data to be displayed (this can be accomplished by the use of appropriate instructions in the content display system scheduling instructions 313 that check for the compatibility of sets of content data with the existing set of application instructions 310) when the attention manager is operating. The content data update instructions 332 can also include instructions that ascertain the current version of the application instructions 310 and, for each set of content data 300 that is incompatible with the current version of the application instructions 310, seek to obtain, at the time scheduled for an update, an updated set of content data 300 that is compatible with the current version of the application instructions 310.

The steps 403 through 406 are advantageous in that they result in the provision of application instructions 310 to a content display system 200 only when such instructions are needed, thus minimizing the number of sets of application instructions that are made available. The steps 403 through 406 also minimize the amount of information that must be transmitted over communication lines to the content display system 200, thereby freeing those lines for other communication and minimizing the cost (i.e., cost of using the communication lines) associated with using the attention manager of the invention.

US 6,788,314 B1

21

Returning to FIG. 4, in the step shown in the block 406 (referred to hereinafter as step 406), the content providing system 203 provides the current set of content data 300 to the content display system 203. (In practice, the set of content data 300 can be provided before, after or simultaneously with provision of the application instructions 310.) Further, as described above, a particular set of content data 300 can exist in different versions that are each compatible with the content display system 203 to which the version of the set of content data 300 is being provided. The step 406 can include a determination as to the version or versions of the set of content data 300 that can be used by the requesting content display system 203, so that a properly formulated set of content data 300 is acquired.

A set of control instructions 320 and content data acquisition instructions 330 (FIGS. 3A through 3C) associated with the set of content data 300 can also be provided, as shown by the step of block 407 (referred to hereinafter as step 407). Typically, a check is made (like that for the application instructions 310 and providing similar benefits) to determine whether the content display system 203 already has a compatible (and/or current) version of the control instructions 320 and/or the content data acquisition instructions 330 associated with the set of content data 300 being obtained.

Each set of content obtained by a content display system 203 can be stored in a database (having any suitable structure) that is stored in a memory of the computer used to implement the content display system 203. The database can also store other information associated with each set of content data 300. This information is discussed in more detail below in the discussion of package files which can be used to convey such information from the content providing systems 203 to the content display systems 203. The package file editor mentioned above can be provided to each content providing system 203 to enable the content provider to easily create a package file for each set of content data 300 provided by that content provider.

Each package file includes a reference to the set of content data 300 (e.g., a network address) to which that package file corresponds. As mentioned above, each package file can also include a variety of other information. For example, the package file can include a specification of the format of the content data 300 (i.e., an indication of the types of content display systems 203 with which the set of content data 300 is compatible) and the type of the content data (e.g., an identification of a particular graphical format, as discussed above). (This information might be specified explicitly or implicitly; alternatively, this information may be passed to the content display system 203 separately from the package file.) The package file can additionally include a text description of the contents of the package file (this could be used, for example, in a user interface that lists descriptions of all of the sets of content data available to a content display system 203 or provided by a content providing system 203). The package file can also include information governing the presentation of the set of content data, such as screen position, special animation effects and display duration (the latter is shown by the View-Time attribute in the Example below). The update information (location and schedule) is also included in the package file. The package file can also include linking information (e.g., network address of an information source) used to implement a link option discussed in more detail below. The content data scheduling information discussed above can also be included in the package file. The package file can also include data structures that can be used to store auditing information, as

22

discussed in more detail below. The package file can also include reference to one or more sets of control instructions 320, each set of control instructions 320 enabling display of the set of content data 300 by a content display system 203 having a particular architecture, or enabling display of clips of particular types.

The following Example illustrates how a package file for use with the invention could be constructed. The package file of this Example does not include all types of information that could be included in a package file; it is to be understood that other types of information (as discussed above, for example) could be included in such a package file, expressed in a similar manner to that shown in the Example. In this illustration, the package file is constructed in an object-oriented manner. Generally, each statement in the package file conforms to the following syntax:

keyword{attribute:value(1) . . . attribute:value(n)}

where "keyword" can be either PACKAGE or CLIP, "attribute" identifies one of the types of information discussed immediately above, and "value" is an identification of particular content for the type of information. There can be any number of "attribute:value" pairs in a statement. In the Example, each attribute:value pair is designated at right by a numeral enclosed in parentheses to aid in the description; this numeral does not form part of the package file shown in the Example.

EXAMPLE

```

PACKAGE {
  Object-Id: 1 (1)
  Object-Type: 1 (2)
  Source: http://www.interval.com/~freiberg/
    Netscreen/Bookreviews/reviews.ass (3)
  Name: Book Reviews: Day 1 (4)
  Description: (5)
  Update-Frequency: 720 (6)
}
CLIP {
  Object-Id: 16919316 (7)
  Object-Type: 2 (8)
  Source: http://www.interval.com/~freiberg/
    Netscreen/Bookreviews/1%20Day%20Book/
    bookreview-1-a1.gif (9)
  Name: Anger (10)
  Description: Book Review (11)
  Update-Frequency: 0 (12)
  View-Time: 15 (13)
  Followup-URL: http://www.randomhouse.com/
    knopf/ (14)
  Linked-To-Following: 1 (15)
}
CLIP {
  Object-Id: 16919384 (16)
  Object-Type: 2 (17)
  Source: http://www.interval.com/~freiberg/
    Netscreen/Bookreviews/1%20Day%20Book/
    bookreview-1-a2.gif (18)
  Name: Emotional Intelligence (19)
  Description: (20)
  Update-Frequency: 0 (21)
  View-Time: 15 (22)
  Followup-URL: http://www.randomhouse.com/
    knopf/ (23)
  Linked-To-Following: 0 (24)
}

```

The first part of line 1 indicates that the following describes a package file. The remainder of line 1 and line 2 are used in debugging and are not relevant to the invention. Line 3 specifies a network address that identifies the location of the package file. (Note that the type of the package file is suggested in line 3 by the file extension .ass, though this

US 6,788,314 B1

23

extension is not necessary to specify the type. This extension can be used to implicitly specify the type of the package file to the content display system 300.) Lines 4 and 5 each give a description of the contents of the package file that can be used, for example, in a user interface to identify the package file. Line 6 specifies the frequency of acquisition of updates to the set of content data 300 and related instructions that are described by the package file. (In this Example, the frequency is specified in minutes.) The first part of line 7 indicates that the following describes a clip in the package file. The remainder of line 7 and line 8 are similar to lines 1 and 2. Line 9 specifies a network address that identifies the location of the clip. (Similarly to line 3, the type of the clip is suggested in line 9 by the file extension .gif, though, again, this extension is not necessary to specify the type.) Lines 10, 11 and 12 are similar to lines 4, 5 and 6. (Note that, in line 12, the specification of "0" for the update frequency indicates that the clip is never updated.) Line 13 specifies the duration of display for this clip. (In this Example, the duration is specified in seconds.) Line 14 specifies a network address of an information source to which a link can be established during display of this clip. (This aspect of the invention is described in more detail below with respect to the "more" option 600d of FIG. 6.) Line 15 specifies the number of additional clips that are part of this package file. Lines 16 through 24 are similar to lines 7 through 15.

When a content display system 200 receives a new package file from a content providing system 202, the content display system 200 first determines whether the contents of a corresponding package file (i.e., either the same package file or an earlier version of the package file) already exist as part of the database. This can be done by scanning a list of entries in the database that each indicate the presence of the contents of a particular package file. If the contents of a corresponding package file are not present, then a new entry is created in the list and the contents of the new package file are stored as part of the database (in accordance with the structure of the database). As part of the process of storing the contents of the package file, the contents are transformed into a form that is compatible with the architecture of the content display system 200 (this is enabled by the installation instructions 313 discussed above).

As described above, the database of content data and related information is constructed from a package file that can have a particular format, as illustrated in the Example above. However, generally, such a database can be constructed from files having any format (e.g., an ASCII file) that enables specification of the information described above that a package file includes.

Returning to FIG. 4, as indicated above, when a set of content data 300 is obtained (step 405), corresponding control instructions 320 and content data acquisition instructions 330 are also obtained (step 407) if such instructions have not already been acquired by the content display system 200. In particular, content data update instructions 333 can be obtained, so that updates to the set of content data 300 and/or the associated control instructions 320 and content data acquisition instructions 330 can be obtained in the future. As mentioned above, the content data update instructions 333 include a description of the location of the content providing system 202 from which the updates can be obtained as well a schedule of times at which such updates should be obtained.

In the step shown in the block 408 (referred to hereinafter as step 408), a determination is made as to whether it is time to update the set of content data 300. The update schedule discussed above is used for this purpose. As long as the

24

schedule indicates that no update need be obtained, the method 400 continues executing the step 408, thereby continuously monitoring whether an update need be obtained. The monitoring of step 408 could be implemented, for example, by a procedure that monitors the content display system computer clock and indicates that an update should be obtained when the clock time is equal to a time in the update schedule.

The update schedule can be established according to any desired criteria. For example, preferably, though not necessarily, the step 408 (and the steps 409 and 410 discussed below and, as necessary, the steps 413 through 417 discussed above) of the method 400 operates at any time that the computer (or computers) with which the content display system 200 is implemented is on, even when the attention manager is not operating. Thus, the update schedule could be established so that updates are obtained during the middle of the night, when charges for communication with content providing systems 202 are cheaper. Preferably, then, at least this part of the content display system 200 is implemented on a computer that is always on, so that such cheap communications time can be utilized for obtaining updates. This can be particularly feasible if the content display system 200 is implemented on a client-server network in which at least the content data update instructions 333 are executed by a server computer which remains on at all times. If, however, the computer on which the content data update instructions 333 are executed is turned off at a time when an update is scheduled to be retrieved, then the update can occur immediately after the next time that the computer is turned on.

This aspect of the content data acquisition instructions 333 can be implemented, for example, using a communications daemon that is part of the content data update instructions 333. When the content data update instructions 333 are acquired by a content display system 200, the daemon is inserted into a startup file that is executed at the beginning of operation of the operating system of the computer with which the content display system 200 is associated. The daemon causes a connection to be made to each location from which the content data update instructions 333 indicate that an update is to be acquired. For example, if the computer uses a Windows operating system, the daemon initiates a WinSock TCP/IP connection to enable connection to be made to the locations of the updated sets of content data 300.

Returning to FIG. 4, once it is determined that an update of the set of content data 300 should be obtained, then, in the step shown in the block 409 (referred to hereinafter as step 409), the location of the appropriate content providing system 202 is ascertained from the scheduling information, and that location is accessed.

In the step shown in the block 410 (referred to hereinafter as step 410), a determination is made as to whether an updated set of content data 300 is available on the content providing system 202. If an updated set of content data 300 is not available, then the step 408 begins executing again, continuing until the update schedule indicates that it is again time to check for an updated set of content data 300. If an updated set of content data 300 is available, then the method 400 returns to the step 408, and an updated set of content data 300 and, if necessary, related control instructions 320 and content data acquisition instructions 330 are provided to the content display system 200 (i.e., an appropriate package file is provided to the content display system 200). As discussed above, the content display system 200 compares the version of the package file contents stored in the database to the contents of the version of the package file being newly provided, and makes changes to the database as necessary.

US 6,788,314 B1

25

FIGS. 5, 5A, 5B and 5C together are a flow chart of a method 500 that implements an attention manager according to another embodiment of the invention. FIG. 5 illustrates the relationship between FIGS. 5A-5C. Like the method 100 (FIG. 1), the method 500 is performed by a content display system 300 according to the invention which can be implemented, for example, using a digital computer that includes a display device and that is programmed to perform the functions of the method 500, as described below. Below, the method 500 is described as implemented on such a digital computer, though the method 500 could be implemented on other apparatus. Steps in the method 500 that are the same as steps in the method 100 are shown by like-numbered blocks. Generally, the method 500 differs from the method 100 in that the method 500 provides a number of control options that enable the user to effect particular types of control of the attention manager. While the method 500 and the associated description below illustrate several control options that can be used with an attention manager according to the invention, it is to be understood that an attention manager according to the invention could include any of a number of other options not shown in FIGS. 5, 5A, 5B, and 5C, or described specifically herein.

The attention manager according to this embodiment of the invention can include any suitable user interface to enable the user to specify a control option. FIG. 6 illustrates a computer display screen 600 including one embodiment of such a user interface. The screen 600 displays, in addition to an image generated from a set of content data 300, a dialog box 601 that includes a list of available control options 602a through 602e. The dialog box 601 can remain on the screen 600 during the entire time that the attention manager is operating. The available control options 602a through 602e shown in the dialog box 601—as well as additional control options that could be, but are not, included in the dialog box 601—are discussed in more detail below.

The manner of selecting an option depends upon the available user input device(s). For example, a keyboard could be used to move a cursor to a desired option, which is then selected using the Enter key. Or, a mouse could be used to move a cursor to a desired option, then clicked to select the option. Or, a touch pen could be used to contact the screen 600 (if the screen 600 is a touch-sensitive screen) at an appropriate location to cause a desired option to be selected. Or, an audio command could be issued to a voice recognition system which causes the desired option to be selected.

One control option that can be used with an attention manager according to the invention enables the user to directly terminate operation of the attention manager. In FIG. 6, this is shown as the "exit" option 602a. In the method 500, this option is implemented using the step 107. As discussed above, selection of the "exit" option 602a causes the primary user interaction to begin again (block 101).

Another control option that can be used with an attention manager according to the invention enables the user to terminate display of the currently displayed set of content data and begin display of the next scheduled set of content data. In FIG. 6, this is shown as the "next" option 602b. In the method 500, this option is implemented by the step shown in the block 301.

Yet another control option that can be used with an attention manager according to the invention enables the user to terminate display of the currently displayed set of content data and begin display of the set of content data

26

displayed immediately prior to the terminated set. In FIG. 6, this is shown as the "back" option 602c. In the method 500, this option is implemented by the steps shown in the blocks 303 and 311.

Still another control option that can be used with an attention manager according to the invention enables the user to terminate display of the currently displayed set of content data and remove that set of content data from the schedule so that the set will not be displayed in the future. This option is not shown in FIG. 6. In the method 500, this option is implemented by the steps shown in the blocks 303 and 312. In a particular embodiment, this option can be implemented so that the set of content data is precluded from being displayed only during the current operation of the attention manager. In another particular embodiment, this option can be implemented so that the set of content data is removed from the content display system 300 entirely, i.e., the set of content data is no longer available for display. In this embodiment, the set of content data could only become available for display again if the user takes affirmative steps to re-obtain the set of content data, as described above with respect to step 401 of the method 400 (FIG. 4).

Another control option that can be used with an attention manager according to the invention enables the user to prevent future display of the currently displayed set of content data until that set of content data has been updated. This option is not shown in FIG. 6. In the method 500, this option is implemented by the steps shown in the blocks 304, 313, 321, 322 and 323 (referred to hereinafter as steps 304, 313, 321, 322 and 323, respectively). If this option is selected in step 304, then an update flag is activated. The update flag can be a designated field associated with a particular set of content data in the database that contains all of the available sets of content data. As shown by step 321, the method 500 identifies, before display of a next set of content data in the schedule, the identity of that next set, and determines (step 322) whether the update flag has been activated for that set of content data. If the update flag has not been activated, then, in step 305, the set of content data is displayed. However, if the update flag has been activated, then, in step 323, a determination is made as to whether the set of content data has been updated since the last time that the set of content data was displayed. This step can be accomplished by checking an update monitor flag that can be a designated field of the database that is associated with the set of content data. If the update monitor flag indicates that the set of content data has been updated since the last time that the set of content data was displayed, then the set of content data is displayed (step 100). Otherwise, the method 500 returns to the step 321 to identify the next set of content data in the schedule.

Yet another control option that can be used with an attention manager according to the invention enables the user to specify a level of satisfaction with the currently displayed content data. This option is not shown in FIG. 6. In the method 500, this option is implemented by the steps shown in the blocks 306, 314 and 315. Depending upon the level of satisfaction indicated in the step 314, the schedule can be modified (step 315) to show the set of content data more, less or at different times than was previously the case. This option can be implemented in any appropriate manner; one way is described immediately below.

The content display system scheduling instructions 313 can include instructions that evaluate a probability function each time that a set of content data in the schedule is presented for display, and either display or not display the set of content data dependent upon the evaluation of the

US 6,788,314 B1

27

probability function. The probability function can include consideration of a variety of factors (e.g., the amount of time that has passed since a particular set of content data has been updated), but for implementation of the instant option, the probability function includes a term np , where n is a constant between 1 and 2, and p is a variable that represents a user's preference for a particular set of content data. Initially, the value of p is 0. Each time that a user indicates a like or dislike for a set of content data (by, for example, selecting an appropriate option in a dialog box such as the dialog box 601), the variable p is incremented or decremented, respectively, by a predetermined amount. The content display system scheduling instructions 343 evaluate a stochastic probability function (e.g., a Gaussian probability function) using the evaluated probability function as an argument. If the result of evaluation of the stochastic probability function is "true", then the set of content data is displayed; if "false", then the set of content data is not displayed. As can be seen, then, initially (i.e., when $p=0$), the user has expressed no like or dislike of a set of content data and the set of content data is displayed or not according to other criteria. Incrementing or decrementing p (i.e., expressing like or dislike for a set of content data) causes the term np to increase or decrease exponentially, thereby increasing or decreasing the likelihood that the set of content data will be displayed.

Still another control option that can be used with an attention manager according to the invention enables the user to establish a link with another information source. In FIG. 6, this is shown as the "more" option 602d. (In "wallpaper" embodiments of the invention, this option can be implemented so that any time the user clicks a mouse—or presses the "Enter" key on a keyboard—when the cursor is within the wallpaper, the link is made to the other information source.) In the method 300, this option is implemented by the steps shown in the blocks 306, 314 and 317. Links can be established to any of a variety of information sources and types of information sources. Typically, the link will be made to an information source that provides information that is related to the content data which was being displayed when the link was established. Upon selection by the user of this control option, the information source is accessed and additional information retrieved for presentation to the user. A link can be made, for example, to any information source that is part of a network which can be accessed by the computer with which the attention manager is being used (though it is not necessary that the link be made through a network). For example, the attention manager can be implemented so that links can be established to locations on the World Wide Web using the appropriate URLs. Such links can be established using any of a variety of Web browser software programs, such as the Navigator software program made by Netscape Communications Corp. Links are enabled by appropriately specifying the location (e.g., a network address) of the information source. The location of an information source (or locations of information sources) can be specified by associating the location with the set of content data, for example, in a package file as described and illustrated above.

As shown in the method 300, the attention manager continues to operate during the time that the link is established and the link is established to an information source from which it is possible to return to the attention manager (see step 317). The presentation of the new information to the user can include an appropriate user interface mechanism that allows the user to request such a return to the attention manager. However, the capacity to return to opera-

28

tion of the attention manager may not always exist. In that event, the step 317 is not part of the method 300; rather, the method 300 terminates after the step 316 and the user operates in the environment of the information source from that point forward. Such termination of the attention manager will frequently be the case where the link is made via a network to an information source.

Another control option that can be used with an attention manager according to the invention enables the user to obtain an overview of all of the content data available for display by the attention manager. This option is not shown in FIG. 6, nor is it implemented in the method 300 of FIGS. 3A and 3B. The overview could be presented textually, pictorially or orally. The overview information can be obtained either via a link to another information location (e.g., the location of the application manager 304) as described above or from a memory associated with the content display system 303, the overview information having been communicated to the content display system 303 when a set of content data was obtained.

Still another control option that can be used with an attention manager according to the invention enables the user to maintain display of the currently displayed set of content data 300 until such display is terminated by the user. This option is not shown in FIG. 6, nor is it implemented in the method 300 of FIGS. 3A and 3B. Upon selection of this option, an appropriate user interface could be made to appear that allows the user to specify termination of the display. After termination of the display, the attention manager resumes normal operation, i.e., the next set of content data 300 is displayed.

The dialog box 601 also includes an additional option, the "cancel" option 602e. Selection of the "cancel" option 602e causes the dialog box 601 to be removed from the screen 600. The dialog box 601 can be made to reappear again using any appropriate technique. For example, the application instructions 340 can include appropriate instructions to cause the dialog box 601 to reappear when the user makes an input to the computer using an input device.

As discussed above (see FIG. 2), usage of the attention manager can be audited using audit instructions 340 (FIGS. 3A and 3C) that can be supplied by the application manager 304 to the content display systems 303, either directly or via the content providing systems 302. The audit instructions 340 can include instructions that cause a content display system 303 to record, as the attention manager is used, particular information (audit information) regarding use of the attention manager (or compute such information from other, more basic information recorded by the attention manager). The audit information can be stored by the content display system 303 in an appropriately structured database. The audit information can include, for example, the identity of each set of content data 300 displayed by the attention manager, the number of times that a set of content data 300 was displayed by the attention manager, the frequency (e.g., number of times per week) that a set of content data 300 was displayed by the attention manager, the times at which a set of content data 300 was displayed by the attention manager, a user-expressed satisfaction level for a particular set of content data 300, and the last set of content data 300 displayed to a user before the user either "passively" (i.e., by making an input to the computer with an input device) or "actively" (i.e., by selecting a control option) terminated operation of the attention manager (of interest, since the user presumably was viewing the display screen when such interaction occurred). The audit instructions can also include instructions that compile and/or

US 6,786,314 B1

29

analyze the audit information in a desired manner. The audit instructions 340 can also include instructions that cause audit information to be transmitted to a remote site (e.g., the application manager 201 or a content providing system 202). These instructions can include scheduling instructions that govern when the audit information is so communicated (e.g., after periodic time intervals), as well as instructions that identify the location (e.g., network address) of the remote site. The transfer of audit information can be accomplished, for example, using a conventional electronic mail mechanism, as known to those skilled in the art. The audit instructions 340 can also include instructions that enable the content display system 203 to display audit information. Additionally, the audit instructions 340 can include instructions that enable the user to disable the audit function entirely, or that enable the user to prevent audit information from being transmitted to the application manager 201 and/or to content providing systems 202. These last instructions could also be accompanied by operating instructions that provide a control option or options to the user, in a manner similar to that described above with respect to FIGS. 5A, 5B and 6, that enable the user to select disablement of the audit function. The audit instructions 340 can also include instructions that cause the database of audit information to be erased at an appropriate time, such as after the audit information has been communicated to a remote site.

Auditing of use of the attention manager can be useful to both users of the attention manager and content providers for a variety of reasons. Such auditing can be used, for example, to illustrate to content providers the value of the attention manager as a tool for disseminating the content provider's information, by showing the content providers how many content data display systems 203 are displaying the content provider's content data. The auditing can also give content providers insight into the interests of computer users, enabling the content providers to better target the information that the content providers provide. The auditing can also indicate to a user the amount and types of the information that the user has been receiving.

Various embodiments of the invention have been described. The descriptions are intended to be illustrative, not limitative. Thus, it will be apparent to one skilled in the art that certain modifications may be made to the invention as described without departing from the scope of the claims set out below. For example, though it is contemplated that an attention manager according to the invention will typically be used to occupy the peripheral attention of a human computer user, generally the attention manager can be used to occupy the attention of any sentient being. For example, the attention manager may be useful in occupying the attention of domesticated animals such as dogs or cats, or providing training (i.e., audio that can be repeated) for a "talking" bird such as a parrot.

We claim:

1. A method for engaging the peripheral attention of a person in the vicinity of a display device, comprising the steps of:

providing one or more sets of content data to a content display system associated with the display device and located entirely in the same physical location as the display device;

providing to the content display system a set of instructions for enabling the content display system to selectively display, in an unobtrusive manner that does not distract a user of the display device or an apparatus associated with the display device from a primary interaction with the display device or apparatus, an image or images generated from a set of content data; and

30

auditing the display of sets of content data by the content display system;

wherein the one or more sets of content data are selected from a plurality of sets of content data, each set being provided by an associated content provider, wherein each associated content provider is located in a different physical location than at least one other content provider and each content provider provides its content data to the content display system independently of each other content provider and without the content data being aggregated at a common physical location remote from the content display system prior to being provided to the content display system, and wherein for each set the respective content provider may provide scheduling instructions tailored to the set of content data to control at least one of the duration, sequencing, and timing of the display of said image or images generated from the set of content data.

2. A method as in claim 1, wherein the display device comprises a television.

3. A computer readable medium encoded with one or more computer programs for enabling engagement of the peripheral attention of a person in the vicinity of a display device, comprising:

instructions for providing one or more sets of content data to a content display system associated with the display device and located entirely in the same physical location as the display device;

instructions for providing to the content display system a set of instructions for enabling the content display system to selectively display, in an unobtrusive manner that does not distract a user of the display device or an apparatus associated with the display device from a primary interaction with the display device or apparatus, an image or images generated from a set of content data; and

instructions for auditing the display of sets of content data by the content display system;

wherein the one or more sets of content data are selected from a plurality of sets of content data, each set being provided by an associated content provider, wherein each associated content provider is located in a different physical location than at least one other content provider and each content provider provides its content data to the content display system independently of each other content provider and without the content data being aggregated at a common physical location remote from the content display system prior to being provided to the content display system, and wherein for each set the respective content provider may provide scheduling instructions tailored to the set of content data to control at least one of the duration, sequencing, and timing of the display of said image or images generated from the set of content data.

4. A computer readable medium as in claim 3, wherein the one or more computer programs enable display of an image or images on a display device comprising a television.

5. A computer readable medium encoded with one or more computer programs for enabling engagement of the peripheral attention of a person in the vicinity of a display device, comprising:

instructions for acquiring a set of content data from a content providing system;

instructions for detecting an idle period of predetermined duration; and

instructions for selectively displaying on the display device, after detection of the idle period and in an

US 6,788,314 B1

31

unobtrusive manner that does not distract a user of the display device or an apparatus associated with the display device from a primary interaction with the display device or apparatus, an image or images generated from the set of content data;

wherein the set of content data is selected from a plurality of sets of content data, each set being provided by an associated content provider, wherein each associated content provider is located in a different physical location than at least one other content provider and each content provider provides its content data to a content display system associated with the and located entirely in the same physical location as the display device independently of each other content provider and without the content data being aggregated at a common physical location remote from the content display system prior to being provided to the content display system, and wherein for each set the respective content provider may provide scheduling instructions tailored to the set of content data to control at least one of the duration, sequencing, and timing of the display of said image or images generated from the set of content data.

6. A computer readable medium as in claim 3, wherein the one or more computer programs enable display of an image or images on a display device comprising a television

7. A content display system for engaging the peripheral attention of a person in the vicinity of a display device located in the same physical location as the content display system, comprising:

data acquisition apparatus that enables acquisition of a set of content data;

display apparatus that effects selective display on the display device, in an unobtrusive manner that does not distract a user of the display device or an apparatus associated with the display device from a primary interaction with the display device or apparatus, of an image or images generated from the set of content data;

user input apparatus that enables selection by a user of one or more control options during the selective display of the image or images generated from the set of content data; and

a system control device that controls aspects of the operation of the system in accordance with a selected control option;

wherein the set of content data is selected from a plurality of sets of content data, each set being provided by an associated content provider, wherein each associated content provider is located in a different physical location than at least one other content provider and each content provider provides its content data to the content display system independently of each other content provider and without the content data being aggregated at a common physical location remote from the content display system prior to being provided to the content display system, and wherein for each set the respective content provider may provide scheduling instructions tailored to the set of content data to control at least one of the duration, sequencing, and timing of the display of said image or images generated from the set of content data.

8. A system as in claim 7, wherein the display device comprises a television.

9. A system as in claim 7, wherein:

a link control option enables the user to establish a link with an information location; and

32

the system control device establishes the link with the information location in response to selection of the link control option.

10. A method for engaging the peripheral attention of a person in the vicinity of a display device, comprising the steps of:

acquiring a set of content data from a content providing system;

selectively displaying on the display device, in an unobtrusive manner that does not distract a user of the display device or an apparatus associated with the display device from a primary interaction with the display device or apparatus, an image or images generated from the set of content data;

enabling selection by a user of one or more control options during the selective display of the image or images generated from the set of content data; and

controlling aspects of the operation of the system in accordance with a selected control option;

wherein the set of content data is selected from a plurality of sets of content data, each set being provided by an associated content provider, wherein each associated content provider is located in a different physical location than at least one other content provider and each content provider provides its content data to a content display system associated with the and located entirely in the same physical location as the display device independently of each other content provider and without the content data being aggregated at a common physical location remote from the content display system prior to being provided to the content display system, and wherein for each set the respective content provider may provide scheduling instructions tailored to the set of content data to control at least one of the duration, sequencing, and timing of the display of said image or images generated from the set of content data.

11. A method as in claim 10, wherein the display device comprises a television.

12. A method as in claim 10, wherein a link control option enables the user to establish a link with a information location, the step of controlling aspects of the operation of the system further comprising the step of establishing the link with the information location in response to selection of the link control option.

13. A computer readable medium encoded with one or more computer programs for enabling engagement of the peripheral attention of a person in the vicinity of a display device, comprising:

instructions for acquiring a set of content data from a content providing system;

instructions for selectively displaying on the display device, in an unobtrusive manner that does not distract a user of the display device or an apparatus associated with the display device from a primary interaction with the display device or apparatus, an image or images generated from the set of content data;

instructions for enabling selection by a user of one or more control options during the selective display of the image or images generated from the set of content data; and

instructions for controlling aspects of the operation of the system in accordance with a selected control option;

wherein the set of content data is selected from a plurality of sets of content data, each set being provided by an

US 6,788,314 B1

33

associated content provider, wherein each associated content provider is located in a different physical location than at least one other content provider and each content provider provides its content data to a content display system associated with the and located entirely in the same physical location as the display device independently of each other content provider and without the content data being aggregated at a common physical location remote from the content display system prior to being provided to the content display system, and wherein for each set the respective content provider may provide scheduling instructions tailored to the set of content data to control at least one

34

of the duration, sequencing, and timing of the display of said image or images generated from the set of content data.

14. A computer readable medium as in claim 13, wherein the one or more computer programs enable display of an image or images on a display device comprising a television.

15. A computer readable medium as in claim 13, wherein a link control option enables the user to establish a link with an information location, the instructions for controlling aspects of the operation of the system further comprising instructions for establishing the link with the information location in response to selection of the link control.

* * * * *

EXHIBIT 5



US 6,757,482 B1

(12) **United States Patent**
Naimark et al.

(16) **Patent No.:** **US 6,757,482 B1**
(45) **Date of Patent:** **Jun. 29, 2004**

(54) **ALERTING USERS TO ITEMS OF CURRENT INTEREST**

(75) **Inventors:** Michael Naimark, San Francisco, CA (US); Amir Burgman, Palo Alto, CA (US); Emily Wall, New York, NY (US); Ignazio Moreano, Brooklyn, NY (US); Ralf Felske, San Francisco, CA (US)

(73) **Assignee:** Interval Research Corporation, Palo Alto, CA (US)

(*) **Notice:** Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by days.

(21) **Appl. No.:** 09/604,428

(22) **Filed:** Sep. 3, 2000

Related U.S. Application Data

(60) Provisional application No. 60/178,627, filed on Jan. 28, 2000.

(51) **Int. Cl.** ⁷ G06F 17/30

(52) **U.S. Cl.** 707/10; 707/2; 707/3; 709/203; 709/218

(58) **Field of Search** 707/1-6, 7, 9-10, 707/101, 104.1; 709/203, 218, 219

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,535,362 A	*	7/1996	Ogawa	707/5
5,724,567 A	*	3/1998	Rote et al.	707/2
5,907,322 A		5/1999	Kelly et al.	345/327
6,078,916 A		6/2000	Collins	707/5
6,098,064 A		8/2000	Piccoli et al.	707/2

6,182,067 B1	1/2001	Pressnell et al.	707/5	
6,208,989 B1	*	3/2001	Dochter et al.	
6,360,261 B1		3/2002	Boyd et al.	709/224
6,385,619 B1	*	5/2002	Eichstaedt et al.	707/104.1
6,460,006 B1		10/2002	Herz	707/10

OTHER PUBLICATIONS

Widyantoro, Ioerger, Yen, "An Adaptive Algorithm for Learning Changes in User Interests", Texas A&M University, College Station, TX 77844-3112, (11-99).

Schwab, Pohl and Korychev, "Learning to Recommend from Positive Evidence", GMD-FIT, MMK, D-53754 Sankt Augustin, Germany, 2000.

Guernsey, Lisa, "Mining the 'Deep Web' with Specialized Drills", Jan. 25, 2001, The New York Times.

eNow-Company Background, <http://www.enow.com/html/comp1.html>, Feb. 7, 2001.

"About Upoc", <http://www.upoc.com/about.jsp>, Feb. 7, 2001.

<http://www.upoc.com/about.jsp>, Sep. 25, 2000.

<http://www.worldwide-top100.com/topsites1/>, Sep. 25, 2000.

<http://my.zdnet.com/myzdnet/alerthelp>, Sep. 25, 2000.

* cited by examiner


Primary Examiner—Charles Rones

(74) **Attorney, Agent, or Firm**—Van Pelt & Yi LLP

(57) **ABSTRACT**

Disseminating to a participant an indication that an item accessible by the participant via a network is of current interest is disclosed. An indication that the item is of current interest is received in real time. The indication is processed. The participant is informed that the item is of current interest.

20 Claims, 14 Drawing Sheets



U.S. Patent

Jan. 29, 2006

Sheet 1 of 14

US 6,757,603 B1




Figure 1

U.S. Patent

Jan. 29, 2004

Sheet 2 of 14

US 6,757,683 B1




Figure 2A




Figure 2B

U.S. Patent

Jun. 29, 2004

Sheet 3 of 14

US 6,797,482 B1




Figure 3

U.S. Patent

Jun. 29, 2004

Sheet 4 of 14

US 6,797,683 B1




Figure 4

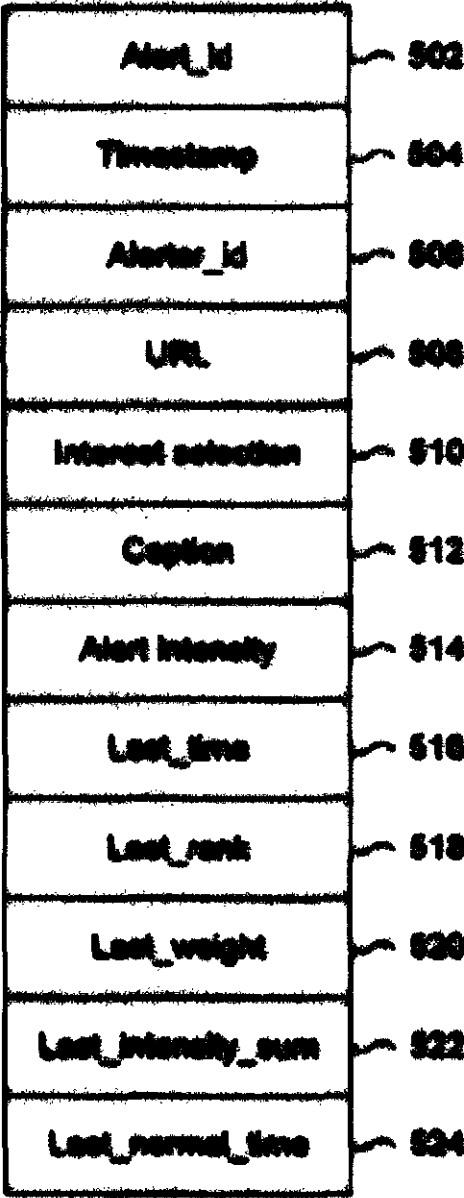


Figure 5

U.S. Patent

Jan. 29, 2004

Sheet 6 of 14

US 6,757,623 B1

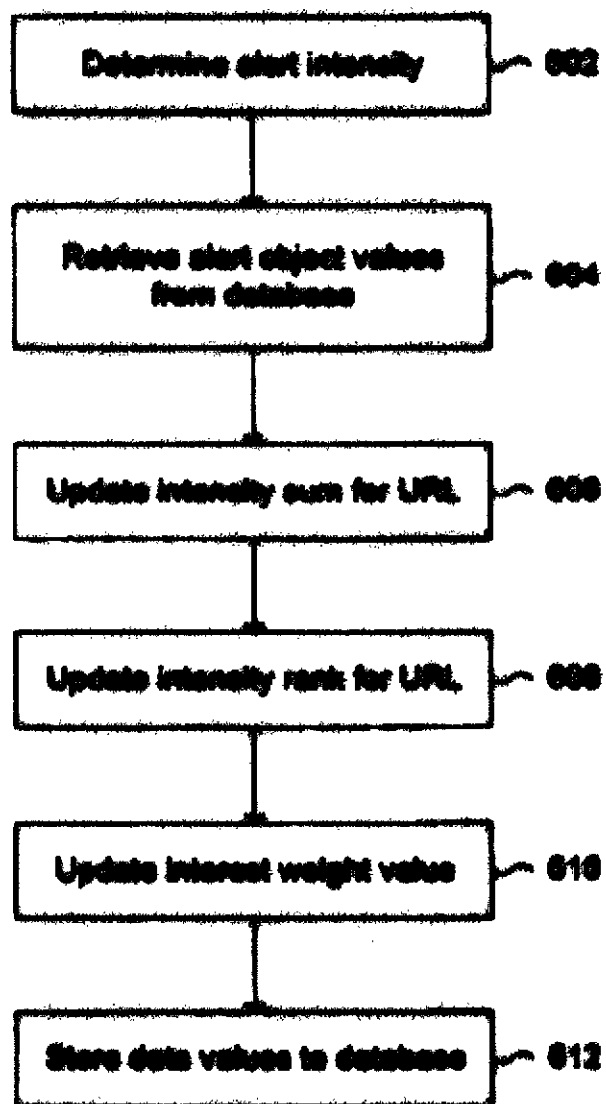


Figure 6

U.S. Patent

Jan. 29, 2004

Sheet 7 of 14

US 6,757,683 B1

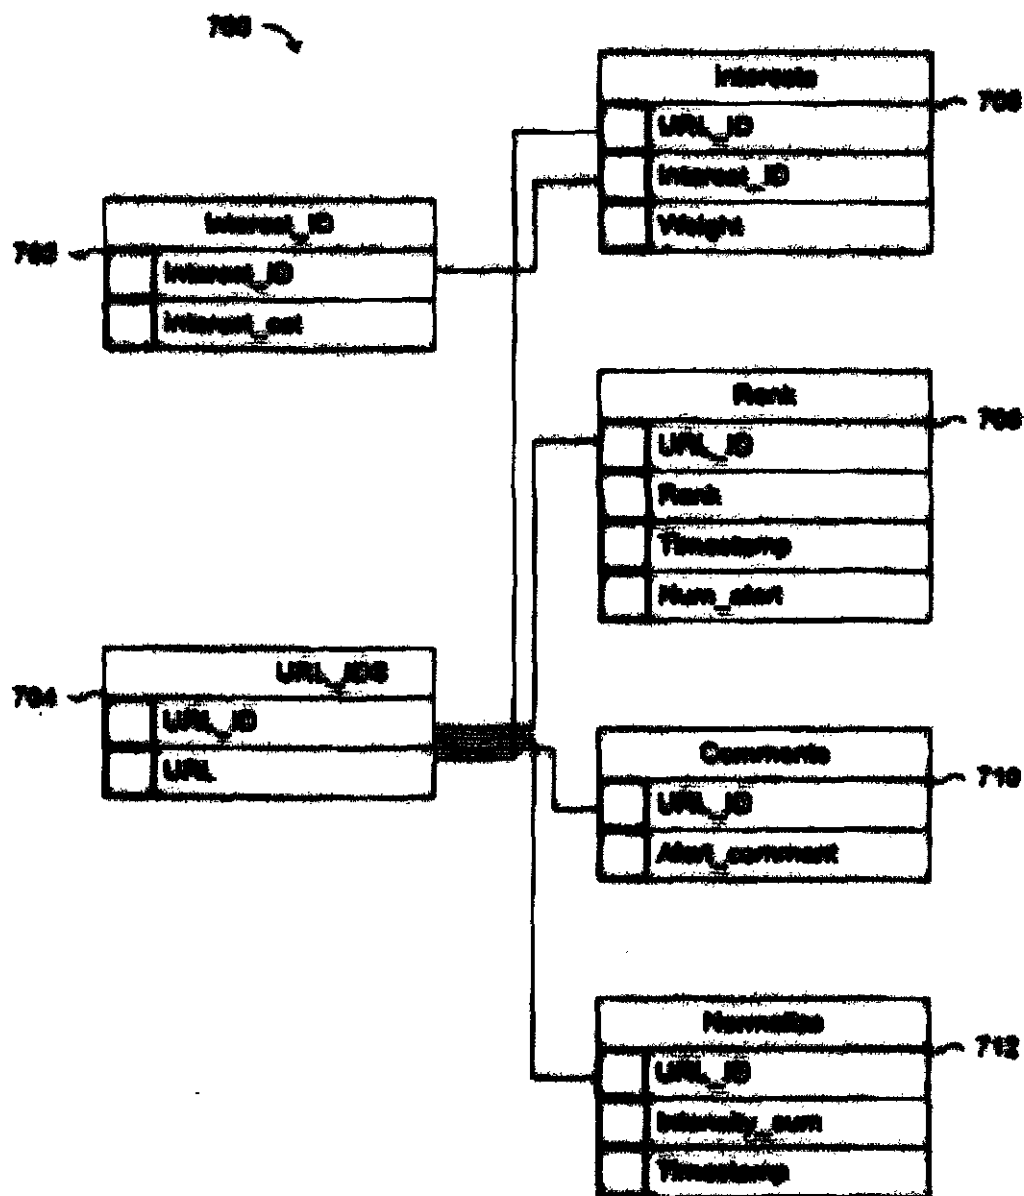


Figure 7

U.S. Patent

Jan. 29, 2004

Sheet 8 of 14

US 6,757,682 B1




Figure 8A

U.S. Patent

Jan. 29, 2004

Sheet 9 of 14

US 6,757,482 B1




Figure 8B

U.S. Patent

Jan. 29, 2004

Sheet 10 of 14

US 6,757,602 B1




Figure 8C

U.S. Patent

Jan. 29, 2004

Sheet 11 of 14

US 6,797,683 B1

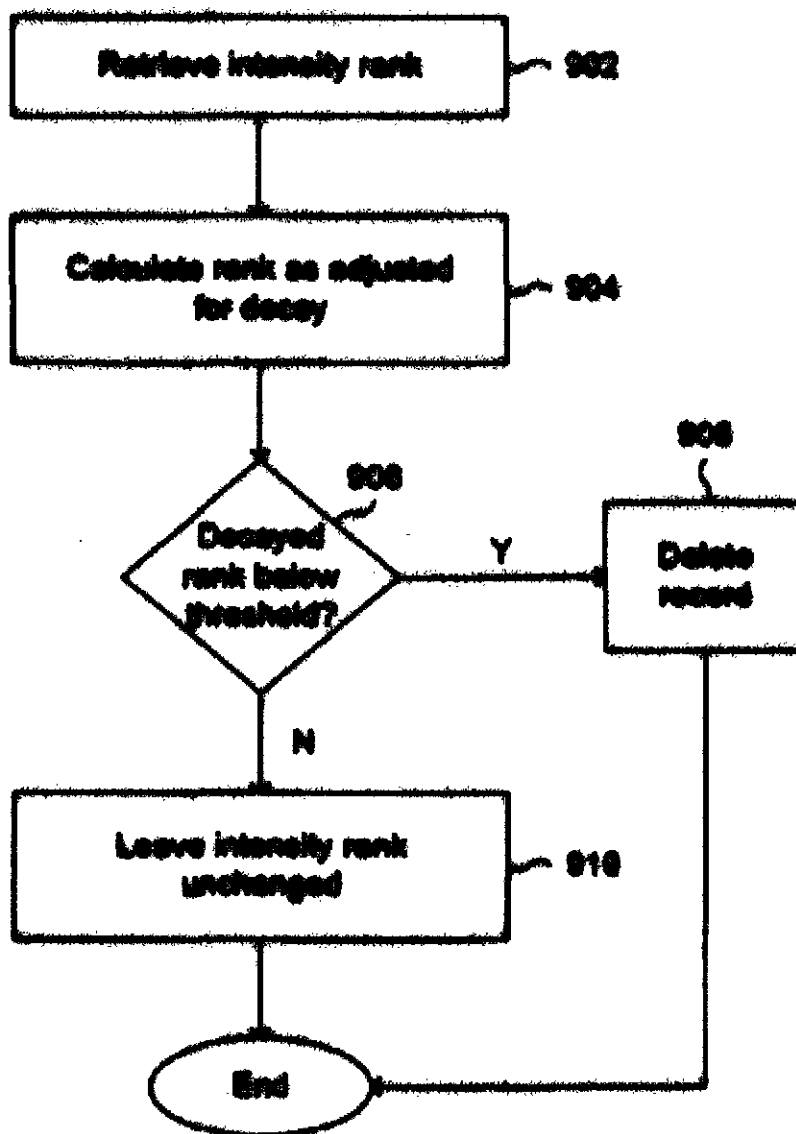


Figure 9

U.S. Patent

Jan. 29, 2004

Sheet 12 of 14

US 6,797,682 B1

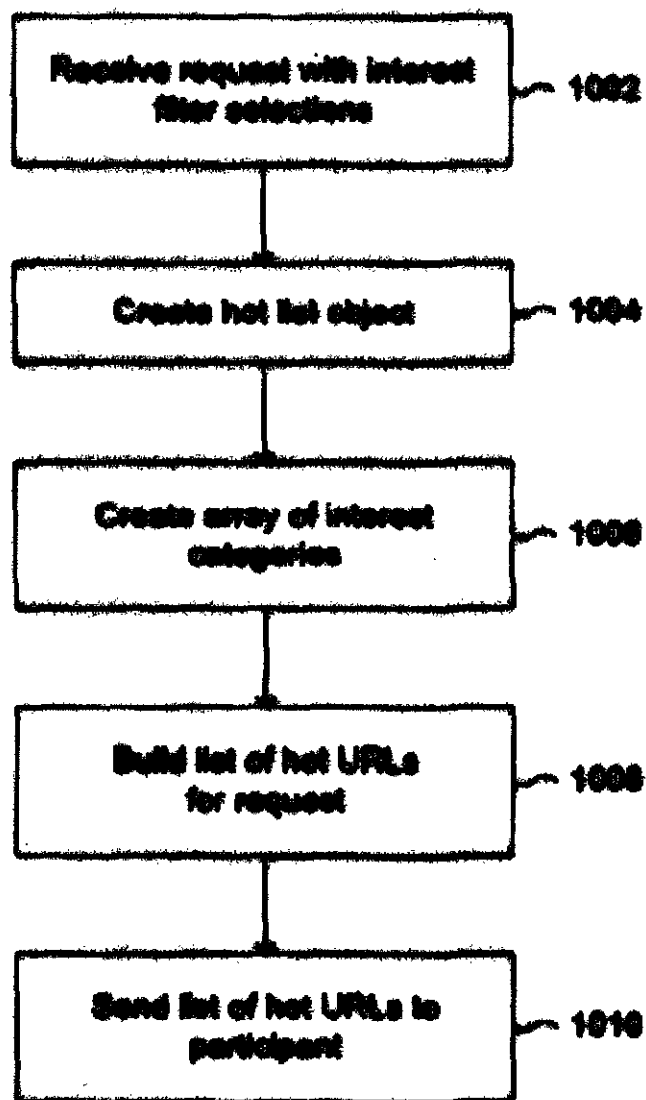


Figure 10

U.S. Patent

Jan. 29, 2004

Sheet 13 of 14

US 6,757,683 B1




Figure 11

U.S. Patent

Jun. 29, 2004

Sheet 14 of 14

US 6,797,603 B1

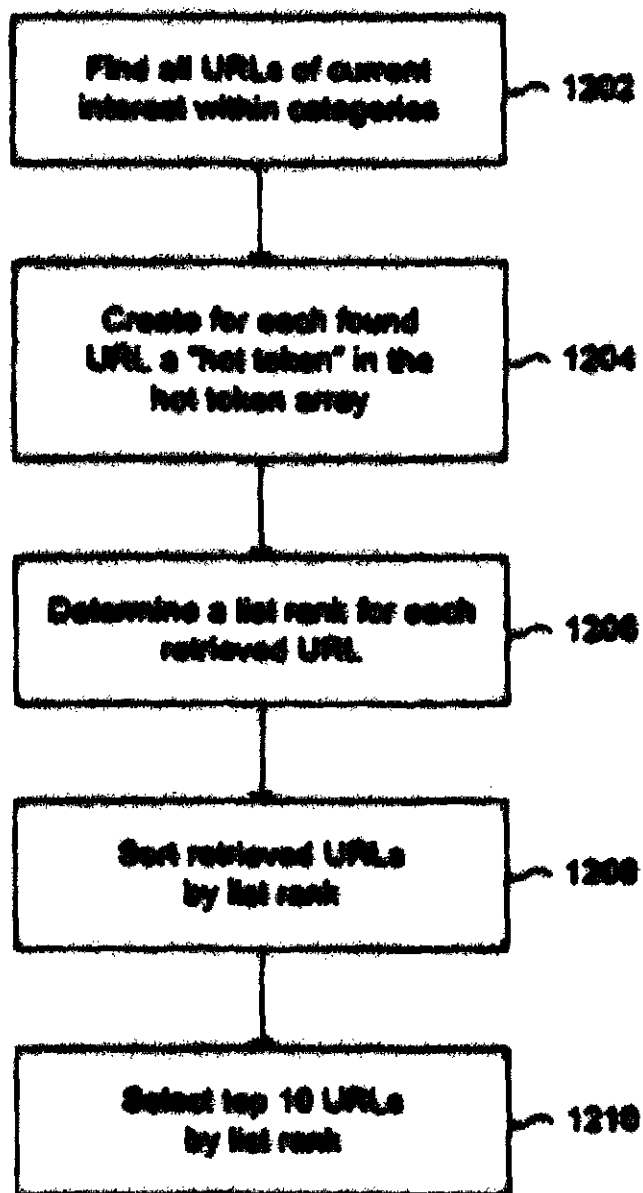


Figure 12

US 6,757,682 B1

1

**ALERTING USERS TO ITEMS OF CURRENT
INTEREST****CROSS REFERENCE TO RELATED
APPLICATIONS**

This application claims priority to U.S. Provisional Patent No. 60/178,627 entitled "Alerting Users To Web Sites of Current Interest And Handling Large Increases In User Traffic" filed Jan. 28, 2006 which is herein incorporated for all purposes.

This application is related to co-pending U.S. patent application Ser. No. 09/656,518, now U.S. Pat. No. 6,356,989 entitled "Quantifying The Level Of Interest Of An Item Of Current Interest" filed concurrently herewith, which is incorporated herein by reference for all purposes; and co-pending U.S. patent application Ser. No. 09/658,346 entitled "Normalizing A Measure Of The Level Of Current Interest Of An Item Accessible Via A Network" filed concurrently herewith, which is incorporated herein by reference for all purposes.

FIELD OF THE INVENTION

The present invention relates generally to communications and computer networks. More specifically, alerting users to dynamic content accessible via a communications or computer network that is of interest at the time of the alert is disclosed.

BACKGROUND OF THE INVENTION

The use of the Internet, and in particular the World Wide Web, and other communication and computer networks has grown dramatically in recent years. The emergence of technologies for broader bandwidth communications, better compression technology, and new and less expensive digital recording and imaging technology, have all contributed to explosive growth in the volume and diversity of content available via communication and/or computer networks, such as the World Wide Web.

However, this proliferation of content, such as audio, image, and video content, presents certain challenges from the perspective of users seeking content of current interest. First, the sheer volume of content available makes it difficult for users to find the content in which they are most interested in accessing at any given time. Apart from having to sort through the enormous volume of content available, much of the content of potentially greatest interest, at least to many users, is dynamic. At certain times, a file or other electronic resource may be of great interest while at other times, or perhaps even most of the time, it is not of great interest or not interesting at all.

For example, thousands of and perhaps in excess of a hundred thousand web cameras, or "webcams", are in use. Webcams are cameras used to provide images of a target of interest via a site on the World Wide Web. Images are updated in varying manners and at varying intervals, depending on the site. A webcam might be used, for example, to provide images of a watering hole in Africa. Typically, users would access a website associated with the webcam to view activity at the watering hole. However, there would be many periods during which nothing of particular interest (e.g., no animals, etc.) would be happening at the watering hole. Conversely, there would be occasional periods when activity of great interest would be occurring, such as the presence of a rare or endangered animal at the watering hole. Users would have no way of

2

knowing when such activity would be occurring, and might miss the most interesting images if they did not happen to check the website at the right time. The same problems arise with respect to files or other electronic resources other than webcam content provided via the World Wide Web, including other media such as audio.

As a result, there is a need for a way to alert users to web content or other electronic resources available via a communications or computer network that are of interest at a particular time. To meet this latter need, there is a need to provide a way to become aware that dynamic web content or an electronic resource other than web content is of interest at a given time, and to quantify the degree or level of current interest. In addition, there is a need to consider the interests of a user when determining which web content or other electronic resources likely will be of the greatest interest to the user.

There is also a need to ensure that interested users receive alerts with respect to web content or other electronic resources that are of interest only to a relatively small community of users, or that are of interest on only relatively rare or infrequent occasions. There is a risk, otherwise, that indications of current interest regarding such files and other electronic resources would be masked by more voluminous or frequent activity with respect to more widely popular or pervasive resources or types of resources (such as pornography sites on the World Wide Web).

SUMMARY OF THE INVENTION

Accordingly, alerting users of items of current interest is disclosed. The level of current interest of a particular file or other electronic resource is determined based on indications received from alerting users. One or more users receive an alert that the item is of current interest. Normalization of the level of current interest of a file or other resource, such as to adjust for items of current interest to a small community or for items of current interest only infrequently, also is described.

It should be appreciated that the present invention can be implemented in numerous ways, including as a process, an apparatus, a system, a device, a method, or a computer readable medium such as a computer readable storage medium or a computer network wherein program instructions are sent over optical or electronic communication links. Several inventive embodiments of the present invention are described below.

Disseminating to a participant an indication that an item accessible by the participant via a network is of current interest is disclosed. In one embodiment, an indication that the item is of current interest is received in real time. The indication is processed. The participant is informed that the item is of current interest.

In one embodiment, a computer is configured to receive in real time an indication that an item is of current interest; process the indication; and inform a participant that the item is of current interest. A database, associated with the computer, is configured to store data relating to the item.

In one embodiment, a computer program product for disseminating to a participant an indication that an item accessible by the participant via a network is of current interest comprises computer instructions for receiving in real time an indication that the item is of current interest; processing the indication; and informing the participant that the item is of current interest.

These and other features and advantages of the present invention will be presented in more detail in the following

US 6,757,682 B1

3

detailed description and the accompanying figures, which illustrate by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

FIG. 1 is a schematic diagram illustrating a system used in one embodiment to alert users to dynamic content of interest at the time of the alert (also referred to herein as an "item of current interest").

FIG. 1A is a series of three screen shots showing three different states of an alert submission display 200 used in one embodiment.

FIG. 2B is an illustration of the data structure used in one embodiment for alerts submitted by an alerting user.

FIG. 3 is a flow chart illustrating a process used in one embodiment to alert users of items of current interest.

FIG. 4 is a flow chart illustrating a process used in one embodiment to receive an alert, as in step 302 of FIG. 3.

FIG. 5 is an illustration of the data structure used in one embodiment for the alert object.

FIG. 6 is a flowchart illustrating a process used in one embodiment to process an alert, as in step 304 of FIG. 3.

FIG. 7 is an illustration of six database tables 700 used in one embodiment to store data concerning alerts received with respect to items of current interest associated with URLs.

FIG. 8A is a flowchart illustrating a process used in one embodiment to update the intensity sum for a URL, as in step 400 of FIG. 6.

FIG. 8B is a flowchart illustrating a process used in one embodiment to update the intensity rank for a URL to reflect the intensity of the current alert.

FIG. 8C is a flowchart illustrating a process used in one embodiment to update the interest category weight for a URL with respect to the interest category indicated in an alert.

FIG. 9 is a flowchart illustrating a process used in one embodiment to purge records for URLs that are determined to be no longer of current interest by calculating a time decayed intensity rank at intervals, even if no new alert has been received, and purging from the database the records for a URL if the time decayed intensity rank is below a prescribed threshold.

FIG. 10 is a flowchart illustrating a process used in one embodiment to disseminate an alert to a participant, as in step 306 of FIG. 3.

FIG. 11 shows an exemplary participant display 1100 used in one embodiment to disseminate alert information to a participant.

FIG. 12 is a flowchart illustrating a process used in one embodiment to build a list of hot URLs responsive to a request, as in step 1000 of FIG. 10.

DETAILED DESCRIPTION

A detailed description of a preferred embodiment of the invention is provided below. While the invention is described in conjunction with that preferred embodiment, it should be understood that the invention is not limited to any one embodiment. On the contrary, the scope of the invention is limited only by the appended claims and the invention

4

encompasses numerous alternatives, modifications and equivalents. For the purpose of example, numerous specific details are set forth in the following description in order to provide a thorough understanding of the present invention. The present invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the present invention is not unnecessarily obscured.

FIG. 1 is a schematic diagram illustrating a system used in one embodiment to alert users to dynamic content of interest at the time of the alert (also referred to herein as an "item of current interest"). The system 100 includes at least one alerting user 102 who accesses dynamic content associated with a uniform resource locator (URL), determines the content is of current interest, and sends an alert indicating that the URL is of current interest, as described more fully below. The system 100 also includes at least one participant 104. In one embodiment, participant 104 provides an indication of the participant's interests and receives a list of URLs providing the location of dynamic content, such as web content on the World Wide Web, that may be of interest to the participant at the time of the alert, as described more fully below. Both the alerting user 102 and the participant 104 are connected to a web server 106 via the Internet. Web server 106 is a computer system configured to present web pages and other web browser readable file, and to receive data from users, via the World Wide Web. Web server 106 is connected to an application server 108 and is configured to provide data to and receive data and instructions from application server 108. Application server 108 is configured to perform the application logic functions described more fully below. In one embodiment, the functions performed by the application server, as described more fully below, are divided among two or more computers so as to optimize the distribution of work load among the computers and to minimize the time the system takes to respond to inputs and queries from users.

When an alert has been received and is being processed, as described more fully below, the application server 108 comprises an alert software object 109 used to store data relating to and perform certain processing with respect to an alert, as described more fully below. The alert software object 109 uses data provided in an alert sent by alerting user 102, along with data retrieved from database 110 associated with the application server 108, to process the alert. Certain of the data that results from the processing performed by alert software object 109 is then stored in database 110. In one embodiment, database 110 is stored in memory in application server 108. In one embodiment, database 110 is stored in a separate structure, such as a database server, connected, either directly or through a communication link, with application server 108.

In one embodiment, when a request from a participant for a list of URLs for items of current interest is received, the application server 108 comprises a hot list software object 112 used to store certain data concerning and perform certain operations with respect to the request from the participant and the response thereto. In one embodiment, the hot list object 112 comprises an interest category array 114. In one embodiment, the interest category array 114 is comprised of one or more interest category objects, each of which stores data relating to one interest category identified in the participant's request as being of interest to the participant. In one embodiment, the hot list object 112 comprises a hot token array 116. The hot token array 116 is

US 6,757,682 B1

5

comprised of a hot token object for each URL of current interest in the database for the category or categories indicated in the participant's request.

As indicated in FIG. 1, an alert sent by an alerting user includes, in one embodiment, at least the URL of the web content considered by the alerting user to be of current interest. In one embodiment an alert may also include an interest selection, meaning a category or subject area to which the alerting user believes the web content relates, and/or a caption in which the alerting user may provide text indicating what the alerting user believes to be of current interest in the web content.

FIG. 2A is a series of three screen shots showing three different states of an alert submission display 200 used in one embodiment. One view is comprised of blank alert submission display 202. Blank alert submission display 202 includes a submission button 204 used to submit an alert with respect to the URL of the web content currently being accessed by the alerting user. Blank alert submission display 202 also includes an interest category selection area 206. In one embodiment, as illustrated in FIG. 2A, the interest category selection area 206 is configured as a pull down menu activated by selecting the downward arrow on the right side of interest category selection area 206. Blank alert submission display 202 also includes a caption area 208 in which an alerting user may enter text associated with the alert, such as text indicating why the alerting user believes the URL to be of current interest. As shown in interest category selection display 212, when the downward arrow button on the right side of interest category selection area 206 is selected, a pull down menu 214 is presented, and an alerting user may select one of the interest categories listed in the pull down menu 214 in the manner well known in the art. As shown in the completed alert submission display 222 of FIG. 2A, the interest category selected by the alerting user is shown in the interest category selection area 206. In the example shown in FIG. 2A, the category selected is "NATURE". In addition, the caption entered by the alerting user, the comment "rhino!" in the example shown in FIG. 2A, appears in the caption area 208 of the alert submission display. As noted above, the alerting party posts the alert to the application server via the Internet and the web server by selecting the submission button 204.

FIG. 2B is an illustration of the data structure used in one embodiment for alerts submitted by an alerting user. The alert includes an ALERTER_ID field 240 in which data identifying the alerting user is provided. The alert also includes a URL field 242 in which the URL of the web content or other electronic resource being accessed by the alerting user when the alert was sent is stored. The alert also includes an INTEREST SELECTION field 244 in which the interest category selected by the alerting user, if any, is provided. Finally, the alert includes a CAPTION field 246 in which the caption entered by the alerting user, if any, is provided.

FIG. 3 is a flow chart illustrating a process used in one embodiment to alert users of items of current interest. The process begins in step 302 in which an alert indicating that an item is of current interest is received. Next, in step 304, the alert is processed. Finally, in step 306, the alert is disseminated to one or more participants, as described more fully below.

FIG. 4 is a flow chart illustrating a process used in one embodiment to receive an alert, as in step 302 of FIG. 3. The process begins with step 402 in which a transmission comprising an alert is received from an alerting user. As

6

noted above, in one embodiment an alert includes at least the URL of the web content being accessed by the alerting user at the time the alert was sent. In one embodiment, as described above, the alert also includes data indicating the identity of the alerting user. In addition, as noted above, the alert may include, at the option of the alerting user, an interest selection and/or a caption for the alert. The process shown in FIG. 4 continues with step 404 in which a new alert software object is created at the application server, such as application server 106 of FIG. 1. Next, in step 406, the data provided in the alert is stored in the alert object. In step 408, a time stamp indicating the time when the alert was received is stored in the alert object. Finally, in step 410, an ALERT_ID, which uniquely identifies the alert and distinguishes the alert and its associated object from other alerts and their associated objects, is obtained and stored in the alert object.

FIG. 5 is an illustration of the data structure used in one embodiment for the alert object. Data field 502 is used to store the ALERT_ID described above. Data field 504 is used to store the time stamp described above. Data fields 506-512 are used to store the ALERTER_ID, URL, INTEREST SELECTION, and CAPTION described above, respectively. ALERT INTENSITY field 514 is used to store a number indicating the intensity or weight to be afforded to the incoming alert. The ALERT INTENSITY is determined as described below. The alert object also stores properties retrieved from various database tables, described more fully below. For example, the alert object includes a LAST_TIME field 516 used to store data retrieved from the database indicating the time of the most recent prior alert. The alert object also includes a LAST_RANK field 518 used to store a numerical ranking retrieved from the database that indicates the overall level or degree of current interest of an item as indicated by all of the alerts that have been submitted with respect to a URL during the current period of activity with respect to the URL through the most recent prior alert. The alert object also includes a LAST_WEIGHT field 520 used to store data retrieved from a database table, as described below, that represents the number of prior alerts received for the URL in the interest category indicated by the current alert, as described more fully below. The alert object also includes a LAST_INTENSITY_SUM field 522 in which the sum of the intensities of all prior alerts for the URL during the current period of activity with respect to the URL, which sum is retrieved from a database table described more fully below, is stored. Finally, the alert object includes a LAST_NORMAL_TIME field 524 used to store the time, retrieved from a database table as described more fully below, when the last normalization calculation was performed.

FIG. 6 is a flowchart illustrating a process used in one embodiment to process an alert, as in step 304 of FIG. 3. The process begins with step 602 in which the intensity of the alert is determined. The term intensity as used herein refers to the weight or value to be assigned to a particular alert regarding an item. In one embodiment, the intensity is a value between 0 and 1. In one embodiment, the value assigned for the intensity is higher if the alerting user selects an interest category for the alert than it would have been if the same alerting party had not selected an interest category. In one embodiment, the intensity value is higher if the alerting party provides a caption for the alert than it would have been if the alerting party had not provided a caption. In one embodiment, the intensity of an alert is increased if it is determined that the alerting party is a party that has provided particularly relevant or helpful alerts in the past, or is trusted for some other reason, such as expertise, academic

US 6,757,682 B1

7

credentials, or reputation within a particular community of interest. In one embodiment, the intensity of an alert is decreased if it is determined that the alerting party has provided unhelpful or erroneous alerts in the past, or if it is determined that the alerting party cannot be trusted as much as other alerting parties for other reasons, such as reputation in the relevant community. In one embodiment, it is possible to provide both an active alert by selecting an alert button and to provide a passive alert by merely accessing a URL with respect to which an alerting party previously submitted an active alert. In one embodiment, an active alert is assigned a higher intensity value than a passive alert.

For example, a passive alert may be arbitrarily assigned a baseline intensity value of 0.3 and an active alert a baseline intensity value of 0.5. For an active alert, 0.1 could be added for each of the following conditions that is satisfied by the alert: an interest category selection was included in the alert; a caption was included in the alert; and/or the source of the alert is particularly trusted. Conversely, 0.1 could be subtracted from the intensity of an alert from a source known to be unreliable. Alternatively, alerts from sources known to be unreliable may be blocked and not assigned any intensity value.

The process illustrated in FIG. 6 continues with step 604 in which data values for the alert object data fields described above that are not included in the alert transmission received from the alerting party are retrieved from the database.

Next, in step 606, the intensity sum for the URL, which is the sum of the intensity values for all of the alerts with respect to the URL, is updated. Next, in step 608, the intensity rank for the URL is updated to reflect the new alert. In step 610, the interest weight value, which represents the number of alerts for a particular URL in which a particular category of interest was indicated, is updated. Finally, in step 612, the updated data values are stored to the database.

FIG. 7 is an illustration of six database tables 700 used in one embodiment to store data concerning alerts received with respect to items of current interest associated with URLs. The database tables 700 include an INTEREST_ID table 702 used to provide a unique identifier, labeled INTEREST_ID in FIG. 7, for each interest category, denominated INTEREST_CAT in FIG. 7. Database tables 700 also include a URL_ID table 704 used to provide a unique identifier, labeled URL_ID in FIG. 7, for each URL.

Database tables 700 also include an INTERESTS table 706 used to store the interest weight, denominated WEIGHT in FIG. 7, for each interest category with respect to which an alert has been submitted for a URL. As noted above, in one embodiment, the weight is the total number of alerts received within a given interest category for a URL. For example, if five alerts indicating the interest category People and three alerts indicating the interest category Nature have been submitted for a URL, there will be two entries for the URL in the interest table, one for each interest category. The weight in the entry for the category People would be "5" and the weight for the URL in the category Nature would be "3".

The database tables 700 also include a RANK table 708 used to store a rank value for each URL associated with an item of current interest, a time stamp when the rank was last calculated, and a data entity denominated NUM_ALERT in FIG. 7, which represents the total number of alerts submitted for the URL.

The database tables 700 also include a COMMENTS table 710 used to store any comment submitted with an alert and to associate each comment with the corresponding URL. Finally, the database table 700 include a NORMALIZE table

8

712 used to store the sum of the intensities of the alerts submitted for a URL (INTENSITY_SUM) and a time stamp indicating when the last normalization was performed.

FIG. 8A is a flowchart illustrating a process used in one embodiment to update the intensity sum for a URL, as in step 606 of FIG. 6. The process begins with step 802 in which the current intensity sum is retrieved from the database, as in step 604 of FIG. 6. If there is no existing record for the URL in the NORMALIZE table (i.e., the alert being processed is the first alert for the URL), a URL_ID is assigned for the URL, a record for the URL is created in the NORMALIZE table, and the retrieved current intensity sum is set to zero. Next, in step 804, the intensity sum is incremented by the amount of the intensity of the current alert. For example, if the previous intensity sum was 4.7 and the intensity for the current alert was 0.5, the intensity sum would be incremented to the value of $4.7+0.5=5.2$. Finally, in step 806, the intensity sum time stamp stored in NORMALIZE table 712 shown in FIG. 7 (which is the same as the LAST_NORMAL_TIME stored in field 834 of FIG. 8) is updated to the time stamp of the current alert. In one embodiment, the intensity sum is updated, and a normalization is performed as described more fully below, each time a new alert is received for a URL. In such an embodiment, the time stamp stored in the NORMALIZE table 712 of FIG. 7 will be the same as the time stamp stored in the RANK table 708 of FIG. 7, as both the rank and the intensity sum are updated each time an alert is received.

FIG. 8B is a flowchart illustrating a process used in one embodiment to update the intensity rank for a URL to reflect the intensity of the current alert. The process begins with step 822 in which the current intensity rank is retrieved from the database, as in step 608 of FIG. 6. As shown in FIG. 7, in one embodiment, this value is retrieved from the RANK table 708. If there is no entry in the RANK table for the URL, i.e., the alert being processed is the first alert for the URL, a record in the RANK table is created for the URL (identified by the URL_ID assigned to the URL) and the current intensity rank is set to zero. Next, in step 824, the intensity rank is updated to reflect the intensity of the current alert. In one embodiment, if the current alert has been received within a predetermined time interval r after the last alert for the URL, the updated intensity rank is a function of the last rank and the intensity of the current alert in accordance with the following formula:

$$r' = (k-r) * I_{alert} + r$$

Where k is the maximum intensity value, which as noted above is one in one embodiment, r is the last rank, r' is the updated rank, and I_{alert} is the intensity value for the current alert. Restating the formula to reflect the fact that in one embodiment, the maximum intensity level $k=1$, the formula becomes:

$$r' = (1-r) * I_{alert} + r$$

If an alert is the first alert received for a URL, the last rank is considered to be zero ($r=0$) and the above formula results in the new rank being equal to the intensity value for the current alert. For example, if the intensity value for the current alert is 0.5, the updated rank $r' = (1-0) * 0.5 + 0 = 0.5$. If a subsequent alert of intensity 0.6 is received, the formula results in the updated intensity rank being calculated as follows:

$$r' = (1-0.5) * 0.6 + 0.5 = 0.8$$

As the example illustrates, so long as additional alerts are received within the time interval each incoming alert will

US 6,757,682 B1

9

cause the intensity rank for the URL to increase until the intensity rank approaches the maximum intensity value k (in the example, the rank would approach $k=1$). The speed with which the intensity rank for a particular URL approaches the maximum value k depends on the intensity value of the incoming alerts and the frequency with which alerts are received.

In one embodiment, if the predetermined time interval τ referred to above has expired between the last alert and the current alert, the updated intensity rank is calculated by a modified formula which reduces the updated intensity rank in accordance with an exponential decay function that effectively adjusts the updated intensity rank downward to account for the passage of time between the last alert and the current alert. All other things being equal, this adjustment would result in a site that received alerts more frequently to have a higher rank than a site that received alerts separated by more than the predetermined time interval. To determine the updated intensity rank as adjusted for the passage of time, the following formula is used in one embodiment:

$$r' = (k - r) * I_{alert} * \tau * e^{-\alpha(\Delta t - \tau)}$$

In this formula, k , τ , and I_{alert} are the same as above, α is the weight assigned to the decay function (a higher value for α will result in a greater amount of decay per unit time), Δt is the amount of time that has elapsed between the current alert and the previous alert, and τ is the predetermined time interval referred to above.

In one embodiment, the updated intensity rank is normalized by multiplying the updated intensity rank by two factors. The first factor is a low frequency enhancement factor designed to enhance the intensity rank of URLs with respect to which alerts are received relatively less frequently relative to the intensity rank of URLs regarding which alerts are received more frequently. The purpose of this enhancement factor is to ensure that sites that are of current interest only from time to time are not masked by the intensity ranking calculated for sites that are of current interest more frequently. In one embodiment, the low frequency enhancement factor is the time of the current alert minus the time of the last update to the intensity rank.

The second factor by which the updated intensity rank is multiplied is a low volume enhancement factor. The purpose of this factor is to ensure that the intensity rank of URLs that are of current interest only to a smaller community of users will not be overshadowed by the intensity rank of URLs that are of current interest to a large community. In one embodiment, the low volume enhancement factor is the inverse of the intensity sum for the URL. Accordingly, in one embodiment, the normalized intensity rank is determined by the following formula:

$$r'' = r' * (I_{current} - I_{sum}) * 1/n$$

Where r'' =normalized intensity rank

r' =updated intensity rank before normalization

$I_{current}$ =timestamp of current alert

I_{sum} =timestamp of first alert for URL

n =intensity sum=sum of all alert intensities for URL.

Once the intensity rank has been updated and normalized, the process shown in FIG. 8B continues with step 836 in which the time stamps for the normalization and intensity rank tables are updated to the time stamp of the current alert.

FIG. 8C is a flowchart illustrating a process used in one embodiment to update the interest category weight for a URL with respect to the interest category indicated in an

10

alert. The process begins with step 843 in which the database is queried to determine if a record exists for the URL for the interest category indicated in the alert. In step 844, it is determined whether the query performed in step 843 identified an existing database table entry for the URL for the interest category indicated in the alert (i.e., whether a prior alert indicated the same interest category for the URL). If it is determined in step 844 that a database entry does not exist for the interest category with respect to the URL, the process proceeds to step 846 in which a record in the INTEREST table is created for the URL with respect to the interest category of the alert. The process then proceeds to step 850 in which the weight value is incremented for the URL with respect to the interest category by increasing the value from zero to one for the new record.

If it is determined in step 844 that there is an existing record for the interest category for the alert with respect to the alert URL, the process proceeds to step 848 in which the weight value stored in the record is retrieved. The process then continues to step 850 in which the retrieved weight is incremented by one to reflect the current alert. For example, if the retrieve weight were 7, the weight would be incremented to 8 in step 850 to reflect the current alert.

FIG. 9 is a flowchart illustrating a process used in one embodiment to purge records for URLs that are determined to be no longer of current interest by calculating a time decayed intensity rank at intervals, even if no new alert has been received, and purging from the database the records for a URL if the time decayed intensity rank is below a prescribed threshold. The process shown in FIG. 9 begins with step 903 in which the intensity rank for a URL is retrieved. In one embodiment, the intensity rank is retrieved and process shown in FIG. 9 is performed, at a predetermined arbitrary time interval τ .

The process shown in FIG. 9 continues with step 904 in which an intensity rank adjusted for time decay is calculated for the URL. In one embodiment, the time decayed intensity rank is determined by the following formula:

$$r_t = r * e^{-\alpha(\Delta t - \tau)}$$

Where

r_t =time decayed intensity rank

α =weight of decay function

Δt =time elapsed since last alert

τ =predetermined time interval referred to above

r =stored intensity rank.

As can be seen from the above formula, the time decayed intensity rank decays exponentially over time if no new alerts are received. If it is determined in step 906 of the process shown in FIG. 9 that the time decayed intensity rank is below the intensity rank threshold, the process proceeds to step 908 in which the record for the URL is deleted. If it is determined in step 906 that the time decayed intensity rank is not below the intensity rank threshold, the process proceeds to step 910 in which the intensity rank as stored in the database is left unchanged.

FIG. 10 is a flowchart illustrating a process used in one embodiment to disseminate an alert to a participant, as in step 306 of FIG. 3. The process begins with step 1003 in which a request containing interest category filter selections made by the participant is received. Next, in step 1004, a hot list software object is created at the application server, as shown in FIG. 1 and described above. Then, in step 1006, an array of interest categories, such as the interest category array 114 described above with respect to FIG. 1, is created within the hot list object. Next, in step 1008, a list of hot

US 6,757,682 B1

11

URLs responsive to the request is built. Finally, in step 1009, the list of hot URLs responsive to the request is sent to the participant.

FIG. 11 shows an exemplary participant display 1100 used in one embodiment to disseminate alert information to a participant. The display 1100 includes a URL entry and display area 1101. The URL for the web content or other electronic resource currently being accessed by the participant is displayed in the URL entry and display area 1101, and the participant may enter the URL for the web content or other electronic resource the participant wishes to access manually in the URL entry and display area 1101, as in the URL or address field for a World Wide Web browser. The display 1100 also includes a content display area 1104 in which the web or other content for the URL listed in URL entry and display area 1101 is displayed. For example, if the URL is the URL of web content accessed via the Internet, the web content associated with the URL will be displayed in URL display area 1104.

The display 1100 also includes an interest category filter selection area 1105 in which interest categories are listed along with a check box for each category listed. The participant selects the check box for each interest category for which the participant would like URLs of current interest to be included in the participant's hot list.

In one embodiment, filter selection area 1105 includes for each category a sensitivity entry area (not shown in FIG. 11) to be used to provide an indication of the participant's degree or level of interest. For example, in one embodiment a participant may enter a whole number from 1 to 5, with 1 indicating the lowest level of sensitivity (e.g., the participant does not want to receive a notification regarding a URL in the category unless a significant number of alerts have been received regarding the URL, or only when the intensity rank for the URL exceeds a predetermined, relatively high threshold) and 5 representing the highest level of sensitivity (e.g., the participant wants to receive a notification even if there has only been one or relatively few alerts concerning a URL, or if one or more alerts have been received but the intensity rank for the URL is relatively low).

In one embodiment, a request is sent to the application server automatically at predetermined intervals. The request contains the interest categories that are in the selected state at the time the request is sent. In one embodiment, the display 1100 includes a submit button (not shown in FIG. 11) that, when selected, causes a request containing the interest categories selected by the participant at the time to be posted to the application server via the Internet.

The display 1100 also includes a hot list display area 1103 in which the hot list of URLs returned by the system to the participant in response to a request is presented. As shown in FIG. 11, in one embodiment, each URL is represented by a hypertext link that, when selected, causes the URL of the listed site to appear in the URL entry and display area 1101 and the content associated with the URL to be displayed in the URL display area 1104.

In one embodiment, the display 1100 is modified to include an alert submission display area such as the alert submission display shown in FIG. 2A. This would permit a participant to send an active alert to the application server if the participant encounters a URL of current interest.

FIG. 12 is a flowchart illustrating a process used in one embodiment to build a list of hot URLs responsive to a request, as in step 1009 of FIG. 10. The process begins with step 1201 in which all URLs of current interest within the categories indicated in the request are found.

Next, in step 1204, a "hot token" object is created in a hot token array within the hot list object for each URL found in

12

step 1201, as described above with respect to hot token array 114 shown in FIG. 1. Each hot token object holds the URL_ID, the WEIGHT for the URL with respect to the interest category indicated in the request, the sum of the WEIGHT values for each category associated with the URL in the database, and the intensity rank (RANK) for the URL.

Next, in step 1206, a list rank is determined for each URL retrieved in response to the request. In one embodiment, a list rank value is calculated for each URL and is used to determine the list rank (or the order in which the responsive URLs will be placed to determine which URLs will be provided). In one embodiment, an initial list rank value is calculated for each URL based on the interest category weight(s) for the URL with respect to the interest category or categories in the request, along with the interest weight for any interest category or categories that are associated with the URL in the database but which are not among the categories indicated in the request. In one embodiment, the initial list rank value "v" of a URL number "n" (v_n) is calculated according to the following formula:

$$v_n = \frac{\sum_{k=1}^m \sqrt{f_k}}{\sum_{k=1}^m \sqrt{f_n}}$$

Where v_n = initial list rank value of URL "n"

f_k = interest weight for URL for each request category

f_n = interest weight for each category associated with URL in database

For example, if at the time of the request there had been ten alerts submitted for a particular URL and three of the alerts were associated with a first category, two with a second category and five with a third category, and if a request were received that included among the request categories the first and third categories, the initial list rank value "v" for URL number "n" calculated in accordance with the above formula would be as follows:

$$v_n = \frac{\sqrt{3} + \sqrt{5}}{\sqrt{3} + \sqrt{2} + \sqrt{5}} = 0.74$$

It should be noted that the use of the square root of the weight for each category tends to give relatively greater effect to the weight of interest categories associated with the URL by a minority of alerting users because using the square root reduces the net effect of the greater weight value associated with interest categories indicated by the majority of alerting users. As with the normalization of the intensity rank described above, this has the effect of giving more visibility to matters of interest to a relatively smaller community.

In an embodiment in which the participant indicates a level of sensitivity with respect to each selected interest category, as described above, the formula for the initial list rank value is modified to take into consideration the sensitivity "s" indicated for each category of interest. In one embodiment, the initial list rank value formula is modified as follows:

$$v_n = \frac{\sum_{k=1}^m \sqrt{s_k + f_k}}{\sum_{k=1}^m \sqrt{s_n + f_n}}$$

Where v_n = initial list rank value of URL "n"

f_k = interest weight for URL for each request category

f_n = interest weight for each category associated with URL in database

65

US 6,737,682 B1

13

s_{m1} =sensitivity indicated for request category "1"

s_{mn} =sensitivity indicated for request category corresponding to interest category "n", if any ($s_{mn}=1$ for interest categories not in request).

For example, in the example described above, assume the participant indicated a sensitivity level of 1 with respect to the first category and 5 with respect to the third category, the initial list rank value would be calculated as follows:

$$v_n = \frac{\sqrt{1 \cdot 3} + \sqrt{5 \cdot 3}}{\sqrt{1 \cdot 3} + \sqrt{1 \cdot 2} + \sqrt{5 \cdot 3}} \approx 0.83$$

(As noted above, the sensitivity level s_{m1} used for the second category, having weight "2" in the denominator, is set at "1" because in the example the participant did not select that category.)

The initial list rank value determined by this calculation (0.83) is greater than the initial list rank value found in the above calculation of an initial list rank value in an embodiment in which sensitivity levels are not assigned or considered (0.74). This illustrates the effect of assigning sensitivity levels. The initial list rank value determined in the second calculation, which takes into account a sensitivity level for each category, is higher than it would have been found to be without regard to sensitivity because the participant indicated a higher sensitivity for one of the categories with respect to which alerts had been received for the URL.

In this way, high-sensitivity users are more likely to become aware of and access a URL with respect to which one or more alerts have been received in a category for which the user has indicated a high sensitivity. If such a high-sensitivity user chose to send alerts of their own with respect to the URL, such activity would increase the intensity rank for the URL (as described above) and would tend to propagate the original alert or alerts to lower-sensitivity users (because the intensity rank is factored into the final list rank used to identify the final list of URLs to be provided to a participant, as described below). If such lower-sensitivity users were to send even more alerts, the original alerts would be further propagated to even lower-sensitivity users, and so on.

In one embodiment, the initial list rank value determined by the interest category weights, as described above, is used along with the intensity rank for the URL to calculate a final list rank value for the URL. In one embodiment, the final list rank value for URL number "n" is calculated in accordance with the following formula:

$$v_{n,final}(\alpha + (1-\alpha)v_n)$$

Where

v_n =final list rank value

r_n =intensity rank for URL

α =weight factor ($0 \leq \alpha \leq 1$)

v_n =initial list rank value

In the above equation, the weight factor α determines the relative weight afforded to the intensity rank for the URL and the initial list rank calculated based on the interest category weights as described above. If the value for α is selected to be 1, the final list rank would be equal to the intensity rank for the URL and the initial list rank would not factor into the final list rank at all. Therefore, a higher weight factor will tend to increase the influence of the intensity rank for the URL and decrease the effect of the initial list rank. Stated another way, a low weight factor tends to give more effect to the extent to which the interest categories associated

14

with the URL in the database match the interest categories indicated in the request from the participant. Conversely, a higher weight factor tends to give greater effect to the overall popularity of the URL as measured by the intensity rank.

Once the list rank for each retrieved URL has been calculated in step 1304, in step 1306 the retrieved URLs are sorted by list rank. Then, in step 1308, the top ten URLs by list rank are selected as the hot list of URLs to be sent to the participant in response to the request. The number ten is an arbitrary number and either a fewer number or greater number of URLs may be included.

Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. It should be noted that there are many alternative ways of implementing both the process and apparatus of the present invention. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

What is claimed is:

1. A system for disseminating to a participant an indication that an item accessible by the participant via a network is of current interest, comprising:

a computer configured to receive in real time from a source other than the participant an indication that the item is of current interest; process the indication; determine an intensity value to be associated with the indication and an intensity weight value, and adjusting the intensity value based on a characteristic for the item provided by the source; and; and inform the participant that the item is of current interest; and

a database, associated with the computer, configured to store data relating to the item.

2. A computer program product for disseminating to a participant an indication that an item accessible by the participant via a network is of current interest, the computer program product being embodied in a computer readable medium and comprising computer instructions for:

receiving in real time from a source other than the participant an indication that the item is of current interest;

processing the indication;

determining an intensity value to be associated with the indication and an intensity weight value, and adjusting the intensity value based on a characteristic for the item provided by the source; and

informing the participant that the item is of current interest.

3. A method of disseminating to a participant an indication that an item accessible by the participant via a network is of current interest, comprising:

receiving in real time from a source other than the participant an indication that the item is of current interest;

processing the indication;

determining an intensity value to be associated with the indication and an intensity weight value, and adjusting the intensity value based on a characteristic for the item provided by the source; and

informing the participant that the item is of current interest.

4. The method of claim 3 wherein processing the indication comprises determining the intensity value for the indi-

US 6,737,682 B1

15

cation based on at least one attribute of the indication, the intensity value representing the weight that will be given to the indication.

9. The method of claim 4 wherein processing the indication further comprises calculating an intensity rank for the item based at least in part on the intensity value of the indication, the intensity rank indicating the level of current interest of the item relative to other items.

10. The method of claim 9, further comprising:

associating the item with a category of interest to which the item relates;

receiving from the participant a selection of one or more categories of interest to the participant;

identifying all items of current interest within the selected categories;

ranking the identified items of current interest; and

sending to the participant a list of items of current interest in rank order, the list including at least one of the identified items of current interest;

wherein the ranking of each item is based, at least in part, on the level of current interest of each item relative to other items as indicated at least in part by the intensity rank.

11. The method of claim 3, further comprising receiving a comment relating to the item.

12. The method of claim 3, further comprising receiving data identifying the source of the indication.

13. The method of claim 3, further comprising associating the item with a category of interest to which the item relates.

14. The method of claim 9, wherein the item is associated with a category of interest identified by the source of the indication of current interest.

15. The method of claim 3, wherein the item is one of a plurality of items of current interest, further comprising:

associating the item with a category of interest to which the item relates;

receiving from the participant a selection of one or more categories of interest to the participant; and

identifying all items of current interest within the selected categories.

16

12. The method of claim 11, further comprising:

ranking the identified items of current interest; and

sending to the participant a list of items of current interest in rank order, the list including at least one of the identified items of current interest.

13. The method of claim 12, wherein the ranking of each item is based, at least in part, on the extent to which the categories selected by the participant match the categories associated with the item.

14. The method of claim 11, further comprising receiving an indication of the participant's sensitivity with respect to each category of interest to the participant, whereby an indication of a relatively low level of sensitivity indicates the participant does not want to be informed that an item is of current interest unless one or more indications have been received that indicate a relatively high level of current interest with respect to an item in the corresponding category and an indication of a relatively high level of sensitivity indicates the participant wants to be informed that an item is of current interest even if only one indication indicating a relatively low level of current interest has been received with respect to an item in the corresponding category.

15. The method of claim 14, further comprising:

ranking the identified items of current interest; and

sending to the participant a ranked list including at least one of the identified items of current interest;

wherein the ranking of each item is based, at least in part, on the sensitivity of the participant with respect to each category associated with the item.

16. The method of claim 3, wherein the item is identified by a Uniform Resource Locator (URL).

17. The method of claim 3, further comprising storing data relating to the indication in a database.

18. The method of claim 3, further comprising determining the weight to be given to the indication.

19. The method of claim 3, wherein the indication is received automatically if a participant accesses the item.

20. The method of claim 3, further comprising providing one or more participants with an interface to send an indication that an item is of current interest.

* * * * *

EXHIBIT 6

AolNews.

News Web
Search News and Topics

SEARCH

Gulf Drilling
Moratorium May Be
Lifted EarlyKim Jong Il's Son
Promoted to GeneralObama Signs Small
Business Bill Into Law

GULF OIL SPILL

Gulf Drilling Moratorium May Be Lifted Early

Updated: 4 hours 1 minute ago

A Print Text Size

RSS News

Laura Parker
Contributor

WASHINGTON (Sept. 27) -- The Interior Department plans to issue two new oil drilling regulations later this week, signaling that a six-month deepwater drilling moratorium for the Gulf of Mexico may be lifted early, a presidential panel was told here Monday.

But Michael Bronwich, director of the newly organized Bureau of Ocean Energy Management, told the Oil Spill Commission that drilling will not immediately resume once the moratorium is lifted.

"You are not going to see drilling going on the next day or even the next week," he said. "I haven't talked to the individual operators. I don't know which are already compliant and which will take time."

The new rules impose a "significant number" of new technical requirements on drilling, he said, but they should be no surprise to the industry. They were outlined in a report issued a month after BP's well blew out on April 20.

William Reilly, co-chairman of the Oil Spill Commission, expressed surprise that the ban -- set to expire Nov. 30 -- has not yet been lifted. He said the 33 rigs affected have been inspected. Nine violations were found, and Reilly said he believes they have been corrected.

He also expressed disappointment in Interior Secretary Ken Salazar's plan to reorganize the Interior Department in order to separate its conflicting roles as oil lease salesman and oil rig inspector.



MORE OIL SPILL NEWS

- Gulf Drilling Moratorium May Be Lifted Early
- BP Claims Cuts From Losses Biggest, Outlier Pays No
- Experts Question BP's Internal Investigation of Deep Oil Spill
- Gulf Operator Could Take 2 Years to Recover
- BP's Well Is Dead, but Challenges Remain
- BP's Blow-Out Gulf Oil Well Finally Declared Dead
- BP Considering Blow-Out Well Forever
- Relief Well Complete: BP Needs Final OK
- Top Experts Defend BP's Early Record
- Report: BP Skipped on North Sea Safety Before Gulf Oil Spill

More Stories >

OIL SPILL COUNTER



Related Articles

From AOL News

BP: Series of Calamities Caused Gulf Oil Spill
19 days ago

2 Ways Obama's New Drilling Moratorium is Different Than the Old One
77 days ago

Judge Blocks Obama's Ban on Offshore Drilling
97 days ago

From Our Partners

BP Oil Spill Created Giant Plume of Crude Underwater in the Gulf
Daily Science - 20 days ago

Alabama Sues BP Over Gulf Oil Spill
Politics Daily - 45 days ago

What the Oil-Rig Cook Saw: Alleged MMS Corruption Was Widespread
Politics Daily - 56 days ago

POWERED BY surphace

EXHIBIT 7

EXHIBIT 8

[My eBay](#)
[Sell](#)
[Community](#)
[Contact us](#)
[Help](#)

[CATEGORIES](#)
[FASHION](#)
[MOTORS](#)
[DAILY DEALS](#)
[CLASSIFIEDS](#)

[Learn more](#)

[Back to Search Results](#) | Listed in category: [Electronics > Televisions](#)

[Add to Watch list](#)

NEW Sony EX5 40 in BRAVIA 1080p LCD HDTV TV KDL40EX500

2-3 Year Warranty! Trusted PowerSeller. Very Fast Ship.

Item condition: **New**

Time left: 5d 07h (Oct 01, 2010 07:00:42 PDT)

[Buy It Now](#)

[Add to Watch list](#)

Shipping: **\$39.99** Standard Shipping See more services | See all details
Estimated delivery within 2-5 business days.

Returns: 14 day money back, buyer pays return shipping | Read details

eBay Buyer Protection
eBay will cover your purchase price plus original shipping.

Seller info
espot.deals (56332) **new**
98.8% Positive feedback

Save this seller
See other items
Visit store: **eSpot Deals**

Other item info

Item number: 120624377443
Item location: **Lanham, Maryland, United States**
Ships to: **United States, Canada**
Payments: **PayPal** See details

[Share](#) [Print](#) [Report item](#)

Check out the most watched

Samsung 23" LCD HDTV 1080p P2370HD ...
\$212.99
3d 18h
 Top-rated seller

Samsung 26" Flat Panel LCD TV HDTV ...
\$289.00
26d 18h
 Top-rated seller

Vizio 37" VLS70H 1080p LCD HDTV
\$414.95
3d 11h
 Top-rated seller

19 inch FAV1 Pink 720p LCD HDTV DVD ...
\$239.95
30d 0h
 Top-rated seller

SAMSUNG UN53C8000 53" 240Hz 3D LED ...
\$2,494.99
26d 15h
 Top-rated seller

EXHIBIT 9

facebook

Search

Home Profile Account

Photos of Alyson

Photo 356 of 1280 Photos of Alyson · Alyson's Profile

Previous Next



In this photo: Barack Obama, Alyson

(photos)

From the album: In-law visits! by Alyson

Added April 14, 2009 · Like · Comment



Write a comment...

Share

Tag This Photo

Report This Photo

Photos of Alyson


My Bachelorette Party! In Town!
Alyson and 2 more friends are tagged in this photo from July 2007.

Friends' Photo Albums





Visitors from Seattle! by Alyson


Sponsored

Create an Ad


Y.T.H.

The upcoming PAKI Band that insists on setting feelings are western Pop/Rock and eastern classical influences fusing deeply

Like · 1,418 people like this.


Marriage is hard.
succeedinlife.org

But divorce is harder. Solving marriage problems is a learned skill. This 12-hour course will end the heartache. Reg. NOW for Jan 2011.


www.usavealot.net SAVE \$
usavealot.net

WOW \$8.99 WOW. SAVE MONEY VISIT US WWW.USAVEALOT.NET BELIEVES IN REAL SAVINGS FOR THE

EXHIBIT 10

Google finance

GOOG

Example: "CSCO" or "Google"

Company

Summary

News

Option chain

Related companies

Historical prices

Financials

Markets

News

Portfolios

Stock screener

Google Domestic Trends

Recent quotes

GOOG	616.72	0.28%
AAPL	308.17	1.31%
APPL	92.00	0.00%
YZ	32.94	1.67%

Create portfolio from quotes

Google Inc. (Public, NASDAQ:GOOG) [Watch this stock](#)**616.72****+1.72 (0.28%)**

Real-time: 11:45AM EDT

NASDAQ real-time data - [Disclaimer](#)


Currency in USD

Range 615.03 - 620.00 P/E 24.99
 52 week 433.63 - 629.51 Dividend -
 Open 618.67 EPS 24.71
 Vol / Avg. 848,952.00/3.58M Shares 318.7MM
 Mkt cap 196.80B Beta 1.21

Compare: Enter ticker here☐ Dow Jones ☐ Nasdaq

1000.2

Screen: 1d 5d 1m 3m 6m YTD 1y 3y 10y All
 Oct 28, 2010 - Nov 02, 2010 -0.88 (-0.14%)

Settings | Plot feeds | Technicals | [Link to this chart](#)

Volume delayed by 15 mins.

Related companies Show: **Most Recent Annual**[Add or remove columns](#)

		Valuation			
Company name		Price	Change	Chg %	Mkt Cap
GOOG	Google Inc.	616.72	+1.72	0.28%	196.80B
MSET	Microsoft Corporation	27.16	+0.21	0.77%	232.53B
YHOO	Yahoo! Inc.	16.22	+0.07	0.44%	21.84B
AOL	AOL Inc.	25.64	+0.12	0.47%	2.74B
AAPL	Apple Inc.	308.17	+3.99	1.31%	282.73B
BIDU	Baidu.com Inc. (ADR)	107.86	+0.71	0.66%	37.52B
ADBE	Adobe Systems Incorporated	28.52	+0.81	1.06%	14.51B
NOK	Nokia Corporation (ADR)	10.62	+0.16	1.59%	39.85B
ORCL	Oracle Corporation	29.34	+0.21	0.72%	147.62B
IBM	Intl. Business Machine...	144.07	+0.75	0.52%	179.11B

Find more results for goog

Dow Jones	11,190.65	0.59%
Nasdaq	2,524.49	0.78%
Technology		0.74%
GOOG	616.72	0.28%

A Zacks Analyst Blog Highlights: Expedia, Procter & Gamble, Priceline, Orbi...
Benzinga - 36 minutes ago

B Analyst Research on Baidu Inc. and Google Inc. - Profits Up for Internet Info...
MarketWatch - 2 hours ago

C 3 Top Tablets to Tempt iPad Buyers
TheStreet.com - 8 hours ago

D Vendors Will Debut Google Chrome Smartbooks Later This Month (GOOG)
Investor's Business Daily - 14 hours ago

E Tank Stock Update: Microsoft Corp.

All news for Google Inc. »

Subscribe

NO

999

TRADE FREE FOR 60 DAYS

E TRADE GET UP TO \$500 +

Events [Add GOOG to my calendar](#)

Oct 14, 2010
Q3 2010 Earnings Conference Call

Oct 14, 2010
Q3 2010 Earnings Release

EXHIBIT 11

Home > Business & Finance > Event Planning

4imprint Guaranteed Low Prices. Expert Service. Ships On Time or It's Free www.4imprint.com

Food Gifts & Baskets Harry & David Specialty Food Gifts. Great Service & Gift Packaging. www.HarryandDavid.com

Gift Basket Ideas Perfect Business Gift Baskets Hand-Picked By Gift Experts. www.Gifts.com

Ad by Google

Q corporate gifts

About.com Event Planning

Event Planning Plan an Event Event Planning Careers

Promotional Products Should Be Practical, Relevant and Creative

What You Should Know When Seeking Creative Corporate Gift Ideas

By Rob Hard, About.com Guide

See More About: [promotional products](#) [corporate gift ideas](#) [creative client gifts](#)

Promotional products are perhaps the most popular corporate gifts. And it probably wouldn't surprise many to learn that it seems that the "pen" is the most relied upon business gift among the hundreds of thousands of choices that exist. Sure, they're practical and useful – and so are a number of other common business and corporate gifts. But there are so many other choices out there.

It's important to keep the promotional products choices simple and convention. At the same time, simple and convention can break the mould and become more creative and artistic – a beautiful garment or print with an artistic touch and professional design will be cherished.

A themed silk screen on a tote bag serves as a practical, creative and unique corporate gift that will be appreciated by those who receive them.

Courtesy of Monsoon Corporate Business Gifts

1.S. Rang, founder of [Monsoon Corporate Business Gifts](#), Montreal, Canada, offered some of the following tips for event planners and others who are searching for creative corporate business gifts that will really make an impression:

1. Give business gifts that people will use.

Free Event Planning Newsletter!

Enter email address

[Discuss](#) in my Forum

Advertisement

Easy to buy. Easy to use.

BUSINESS AND EMPLOYEE GIFT

from Sodexo

Learn More >>

www.SodexoGifts.com [Ad by Google](#)

Sponsored Links

Unique Personalized Gifts
Free Personalization on Every Gift & Fast 1-2 Day Shipping. Shop Now!
www.PersonalizationMall.com

Cool Gifts for Guys
Gadgets, Barware, Cufflinks & More. Available Only at RedEnvelope
www.RedEnvelope.com/gifts-for-him

Pocket Hand Sanitizer Pen
Imprinted with Your Logo Discount Pricing / Fast Service
www.GreenMonsterPromos.com

Video

Set Up a Mail Merge in Outlook

EXHIBIT 12

[Browse DVD](#)

[Movies, TV shows, actors, directors, genres](#)

[Home](#)
[Genres](#)
[New Releases](#)
[Netflix Top 100](#)
[Critics' Picks](#)
[Award Winners](#)

Dante's Inferno

2010 **DR** 80 minutes

Traveling beyond death to rescue his love, Beatrice, from evil Lucifer, Dante enters the nine circles of hell, battling the vicious demons and monsters he encounters in this spectacular animated retelling of Dante Alighieri's classic morality tale. A companion to EA's video game, the action-packed medieval journey features the voices of Mark Hamill, Victoria Tennant, Vanessa Branch and Graham McTavish.

Cast: Graham McTavish, Vanessa Branch, Steve Blum, Peter Jessop, Mark Hamill, Victoria Tennant, Bart McCarthy, Kevin Michael Richardson, John Paul Karliak, Tom Tate

Director: Mike Disa, Shuko Murase, Victor Cook, Jong-Sik Nam, Lee Seung-Gyu, Sang-Jin Kim

Genres: Anime & Animation, Animation for Grown-ups, Anime Horror, Anime Feature Films, Blu-ray

Format: DVD and Blu-ray

☐ Not Interested

Our best guess for Jay: 2.9 stars
Average of 30,510 ratings: 3.3 stars

[Add to DVD Queue](#)

More like Dante's Inferno

Tench the Movie 2: Daughter of...

Add

☐ Not Interested

Dead Space: Downfall

Play Add

☐ Not Interested

Robotech: The Shadow Chronicles

Play Add

☐ Not Interested

Hellboy: Animated: Blood & Iron

Play Add

☐ Not Interested

Lady Death


Play Add

☐ Not Interested

Some content copyright 1981-2010 Muze Inc. For personal non-commercial use only. All rights reserved.

EXHIBIT 13

EXHIBIT 14



OfficeMax
WORK WITH US™

Welcome RajeshM | Sign Out | My Account | Gift Cards |

Order By Number | Waddy Ad | Store Locator | Contact Us | Track an Order |

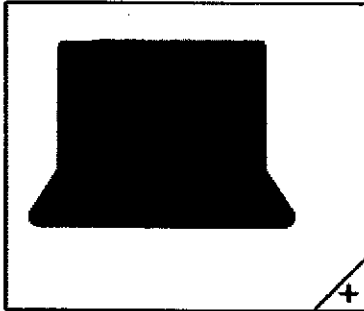
My Cart: \$2,388.96 (4 items)

OFFICE SUPPLIES | PAPER | TECHNOLOGY | INK & TONER | OFFICE FURNITURE

Keyword or Item #

HOME > TECHNOLOGY > COMPUTERS > LAPTOP COMPUTERS AND TABLET PCs

☐ Print ☐ Email a Friend



COMPAQ PRESARIO CQ62-220US NOTEBOOK PC

Avg. Customer Rating: ★★★★★ 4.7 out of 5 (3 reviews)
Read reviews | Write a review

Reg: \$629.99
Save 19% \$500.00
Your Price: \$429.99

INSTANT SAVINGS

Ships in 1 to 3 days
Item #: 22334243
Manufacturer #: CQ62-220US

Featuring: AMD Athlon II Processor, 320 GB Hard Drive, 3GB Memory, Windows 7 Home Premium, 15.6" diagonal widescreen

Features

- VISION Technology with an AMD Athlon II Dual-Core Processor for Notebook PCs P320 (2.1GHz, 1MB L2 cache)
- Genuine Windows® 7 Home Premium (64-bit)
- 15.6" diagonal high-definition HP BrightView widescreen LED display (1366 x 768)
- 3GB DDR3 system memory
- 320GB 5400rpm SATA hard drive
- LightScribe SuperMulti 8X DVD±R/RW drive with Double Layer support
- 802.11b/g/n WLAN
- Textured HP Imprint finish in matte black
- 6-cell Li-Ion battery
- 1-year limited warranty

Quantity:

ADD TO FAVORITES

TECHNICAL DETAILS **ACCESSORIES** **CUSTOMER REVIEWS**


The Compaq Presario CQ62 series offers the essential features and applications you need to enhance your productivity. Convenient one-touch keys provide instant access to your e-mail, the Web, and other functions you use most. Get fast, comfortable performance with the CQ62.

Warranty: 3-year limited hardware warranty, 3-year free hardware technical support, 30-day free limited software support

Product Specs:


See Windows 7 in action

RELATED PRODUCTS



Microsoft Office Home & Student 2010
Item #: 21772288
Unit: EA
Ships in: 1 to 3 days
Price: \$149.99


Quantity:



HP Photosmart Premium All-in-One
★★★★★ 3.8 out of 5
Item #: 21987227
Unit: EA
Ships in: 1 to 3 days
Price: Reg: \$499.99 \$149.99

INSTANT SAVINGS

Quantity:



Microsoft Office Home & Student 2010
Product Key Card
Item #: 22311231
Unit: EA

EXHIBIT 15

that was easy:

Find / Products

Home | Favorites Lists | Easy Reorder | Order by catalog item #

Store Locator | Help | Log in

0 items \$0.00 Checkout

Welcome, Log in please. | Your prescription for illness prevention starts here.

Buy Online & ship to store **Fast and FREE DELIVERY** on all orders over \$50

Weekly Ad
Valid 9/28 - 10/02

Back to: [All Digital Cameras](#) / [Kodak EasyShare M530 Digital Camera Series](#)

Share

1 of 8

Special Buy

Reg: \$129.99
Save: -\$12.00
Now: \$117.99
Instant Savings: -\$18.00
Rebate: See Details

Price before rebate: \$99.99
Each

Qty:

Item qualifies for free delivery
Expected Delivery 2 Business Days
Item qualifies for free delivery to a store near you. [Learn more](#)
Check in Store Availability

Item: 635724
Model: 1286663

[Live Customer Support](#)

Product Details

Customer Rating
★★★★★
(14 reviews)
[Write a Review](#)
[Read 14 Reviews](#)

Q&A
[Ask a Question](#), [Share comments](#)
[3 Questions 1 Answers](#)
[Ask a Question](#)

[ShareThis](#)

Sharing your life just got simpler. Just press Share.

Picture this-you're letting loose at a trendy hot spot with friends, snapping pictures so amazing you'll want to share them online as soon as you get home. Any other camera would slow you down with tags, sorts, and uploads. But not the new KODAK EASYSHAREM530 Digital Camera. Kodak's Share button allows you to tag pictures directly from your camera for e-mailing, or uploading to YouTube, FACEBOOK, FLICKR and KODAK Gallery sites, plus e-mail too. The real Kodak Moment happens when you share.

- 12.2 megapixels
- 3x optical/5x digital zoom
- 2.7" indoor/outdoor color LCD display
- Just press Share - One-button upload to send (pictures/video) via email or upload to YOUTUBE, FACEBOOK, FLICKR, or KODAK Gallery
- Reduce blur caused by subject movement or fast-action situations
- Kodak's innovative Smart Capture feature automatically identifies the scene and adjusts camera settings for a perfect picture in just about any environment

Customers who viewed this camera also bought:

	Qty.
<p>SanDisk® 4GB SDHC™ Card ★★★★★ (388 reviews) \$27.99 Each</p>	<input type="text"/>
<p>Kodak Digital Camera Hard Case, Red \$12.99 Each</p>	<input type="text"/>

EXHIBIT 16